



AGP Design Guide

Covering 1X, 2X, 4X and 8X Modes and 0.8 Volt, 1.5 Volt and 3.3
Volt Signaling

Revision 2.0
Intel Corporation
October 2002



THIS SPECIFICATION IS PROVIDED “AS IS” WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification.

§ Other names and brands may be claimed as the property of others.

Copyright © 1996-2002, Intel Corporation. All rights reserved.

Contents

Preface	13
Document Structure and Outline	13
Content	13
Feedback	13
 1 AGP Electrical Design Considerations	15
1.1 AGP Overview.....	15
1.1.1 System Timing Overview	16
1.2 Transmitter.....	21
1.2.1 Internal Clock Generation and Distribution.....	21
1.2.2 Output Buffer Considerations	23
1.2.2.1 Buffer Architecture	23
1.2.2.2 Buffer Design.....	25
1.2.3 Simultaneous Switching Output (SSO) Noise and Pattern Effects.....	30
1.2.4 Transmitter Floorplan	32
1.3 Interconnect.....	32
1.4 Receiver.....	33
1.4.1 Input Receiver Considerations – AGP 2X and AGP 4X	34
1.4.1.1 Input Receiver Considerations – AGP 8X.....	36
1.4.2 Termination on Unused Signals (Pull-ups / Pull-downs): AGP 2X/AGP 4X.....	37
1.4.2.1 Termination of Unused Signals (Pull-ups / Pull-downs) – AGP8X	39
1.4.3 Signal Overshoots and Undershoots – AGP2X/AGP4X.....	39
1.4.3.1 Signal Overshoots and Undershoots – AGP8X.....	39
1.4.4 Voltage / Loading Characteristics of Vref	39
1.4.4.1 Vref Characteristics for AGP 2X Mode.....	39
1.4.4.2 Vref Characteristics for AGP 4X Mode.....	40
1.4.4.3 Vref Characteristics for AGP 3.0 Mode.....	41
1.4.4.4 AGP3.0 VREF Generation Using Detect Signal Circuits	42
1.4.4.5 Vref Margining.....	43
1.5 AGP Vddq Power Delivery	44
1.5.1 Vddq Voltage Sources	44
1.5.2 Motherboard Parasitics	45
1.5.3 Connector Parasitics	46
1.5.4 Power Plane Decoupling	46
1.5.5 Card Parasitics	48
1.5.6 Transient Current Load Waveform.....	49
1.5.7 Connector AC Signal Decoupling Requirements	49
1.6 Motherboard and Expansion Card PCB Layout Considerations.....	50
1.6.1 Motherboard Layout Recommendations	50
1.6.2 Expansion Card Layout Recommendations.....	53
1.6.3 Board Impedance	54
1.6.4 Pull-ups and Termination.....	54
1.6.5 Board Routing Recommendations.....	54
1.6.5.1 Maximum Trace Length Requirements	54

1.6.5.2	Trace Length Mismatch Requirements	56
1.6.5.3	Strobe Trace Routing Considerations.....	56
1.6.5.4	Clock Routing and Skew	56
1.6.6	Interface Signaling Requirements.....	57
1.6.7	IDSEL Routing for Add-in Cards.....	58
1.7	IC Packaging Considerations	58
1.7.1	Signal Skew	59
1.7.2	Signal Loading	59
1.7.3	Signal Routing and Crosstalk	59
1.7.4	Power Pins in the Signal Ball Array.....	60
2	AGP Signals Sensitivity Analysis and Measurement Techniques.....	62
2.1	2x Mode Board Design Guidelines	62
2.1.1	2X Mode Routing Rules.....	62
2.1.1.1	AGP 2X Mode Design Considerations.....	63
2.1.1.2	Source Synchronous Recommendations for Add-in Cards.....	64
2.1.1.3	Source Synchronous Recommendations for Motherboards.....	64
2.1.1.4	Control Signal and Clock Recommendations	65
2.1.2	Simulation Techniques	65
2.1.2.1	Methodology.....	66
2.1.2.2	Buffer Models	67
2.1.2.3	Interconnect Models and Crosstalk Effects	67
2.1.2.4	Simulation Setup	68
2.1.2.5	Simulation Parameters.....	68
2.1.2.6	Flight Time	69
2.1.2.7	Flight Time Skew	70
2.1.2.8	Signal Quality	71
2.1.2.9	Summary of Simulations Required	72
2.1.3	Simulation Results.....	74
2.1.3.1	Flight Time	74
2.1.3.2	Flight Time Skew	74
2.1.3.3	Signal Quality	76
2.2	4X Mode Board Design Guidelines	76
2.2.1	4X Mode Routing Rules.....	76
2.2.1.1	AGP 4X Mode Design Consideration	76
2.2.1.2	Source Synchronous Rules for Motherboards.....	78
2.2.1.3	Source Synchronous Rules for Add-in Cards.....	79
2.2.1.4	Control Signal and Clock Recommendations	79
2.2.1.5	Inter-Symbol Interference	80
2.2.1.6	Board Constraints.....	81
2.2.2	Simulation Techniques	81
2.2.2.1	Methodology.....	82
2.2.2.2	Simulation Parameters.....	84
2.2.2.3	Buffers	84
2.2.2.4	Interconnect.....	84
2.2.2.5	Crosstalk	85
2.2.2.6	Flight Time	87
2.2.2.7	Flight Time Skew	87
2.2.2.8	Signal Quality	88
2.2.3	Simulation Results.....	89
2.2.3.1	Signal Quality Simulations.....	89

2.2.3.2	Minimum Flight Time Simulations.....	89
2.2.3.3	Maximum Flight Time Simulations.....	90
2.2.3.4	4X Transfer Mode Simulations.....	91
2.2.4	Design Recommendations.....	97
2.2.4.1	Board Materials	97
2.2.4.2	Routing Measurement Points	98
2.2.4.3	Escaping Components	98
2.2.5	Measurements	99
2.2.5.1	Ringback.....	99
2.2.5.2	Skew.....	99
2.3	8x Mode Board Design Guidelines	100
2.3.1	8X Mode Routing Rules.....	100
2.3.1.1	AGP 8X Mode Design Considerations.....	101
2.3.1.2	Board Constraints.....	102
2.3.1.3	Source Synchronous Rules for Motherboards.....	103
2.3.1.4	Source Synchronous Rules for Add-in Cards.....	104
2.3.1.5	Control Signal and Clock Recommendations	105
2.3.1.6	Routing Measurement Points	105
2.3.1.7	Package Trace Matching.....	106
2.3.1.8	Ground Plane Reference	106
2.3.2	Simulation Techniques	107
2.3.2.1	Methodology.....	108
2.3.2.2	Simulation Parameters	110
2.3.2.3	Buffers	110
2.3.2.4	Packages	110
2.3.2.5	Connectors.....	110
2.3.2.6	Printed Circuit Board	111
2.3.2.7	Crosstalk	111
2.3.2.8	Sweep Matrix.....	113
2.3.2.9	Skew Calculations.....	114
2.3.2.10	Signal Quality	116
2.3.3	Simulation Results.....	120
2.3.3.1	Signal Quality Simulations.....	120
2.3.3.2	Timing Skew Simulations.....	122
2.3.4	Additional Design Recommendations.....	123
2.3.4.1	Board Materials	123
2.3.4.2	VDDQ Plane	124
2.3.4.3	Connector Impedance Matching.....	124
2.3.4.4	Package Pinout	124
2.3.4.5	Trace Bends and Serpentes	124
2.3.4.6	Measurements	124

3 Manufacturing Test and System Validation Guidelines.....127

3.1	Manufacturing Test Overview	127
3.1.1	Test Platform limitations with Edge Placement Accuracy.....	127
3.1.2	Accurately Assessing Source Synchronous Timings	130
3.1.3	Measuring Pad timings at a Pin	132
3.1.4	Testing with Low Data Rate Test Equipment.....	135
3.2	System Validation Measurement Guidelines.....	137
3.2.1	Measurement Techniques	137
3.2.2	Test Equipment Required.....	137

3.2.2.1	Oscilloscopes	137
3.2.2.2	Scope Probes	137
3.2.2.3	De-Embedding Test Equipment Bandwidth.....	138
3.3	Time Domain Measurements	138
3.3.1	66 MHz clock jitter and skew.....	138
3.3.2	Common Clock Data Timing	138
3.3.3	Source Synchronous Timing	139
3.4	Frequency Domain Measurements	140
3.4.1	VNA Measurements.....	140
3.4.2	Interpreting and Using the Data.....	140
3.5	Waveform Measurements and Parameters	141
3.5.1	66 MHz Clock Jitter and Skew	141
3.5.2	Data to 66 MHz Clock Timing and Voltage Parameters	142
3.5.3	Source Synchronous Timing and Voltage Parameters	142
3.6	Voltage Margin Measurements	143
3.6.1	Data Test Patterns.....	143
3.7	Logical Measurements	143
3.7.1	AGP8X Logic Analyzer Interface	143
3.7.2	LAI Connector.....	146
3.7.2.1	Recommended Probing Solution	147
4	AGP Thermal Design Guidelines	150
4.1	Thermal Overview.....	150
4.2	Thermal Design Recommendations	150
4.2.1	Recommendations for Add-In Card Designers	150
4.2.2	Power Budgets	151
4.3	Recommendations for OEM System Designers	151
4.4	Thermal Performance Recommendations.....	151
4.4.1	Add-in Card	151
4.4.2	Fan Direction.....	152
4.4.3	Venting.....	153
4.4.4	Internal Air Flow	154
4.5	Thermal Design Examples	154
4.5.1	Reference Documents	157
A.	Crosstalk Modeling Techniques	159
A.1	Objective	159
A.2	Flight Time and Skew as a Function of Odd and Even Propagation Modes.....	159
A.2.1	Definitions	159
A.2.2	Flight Time Analysis of Two Conductors Using a Single Conductor Model.....	161
A.2.3	Flight Time/Skew Analysis for More Than Two Conductors	163

A.2.4	Methodology.....	163
A.2.4.1	Example 1.....	163
A.2.4.2	Example 2.....	164
A.2.4.3	Example 3.....	165
B.	Models.....	167
C.	Sensitivity Analysis	173
C.1	AGP-4X.....	173

Figures

Figure 1: System Overview.....	16
Figure 2: AGP Design Components	16
Figure 3: AGP Time Domains.....	17
Figure 4: Clock Design.....	22
Figure 5: Clocking Methodologies	23
Figure 6: Clock Design for AGP 8X/4X/2X System.....	23
Figure 7: Buffer Configurations.....	24
Figure 8: Data and Strobe Buffers	24
Figure 9: Digital Impedance Compensation	26
Figure 10: Binary Encoding.....	26
Figure 11: Increasingly Linear Buffer Curves	27
Figure 12: AGP 4X Pull Down/Pull Up V/I Curve	28
Figure 13: Linearity measurement.....	29
Figure 14: Output Buffer Reference Network	30
Figure 15: Example PCB Layer Stackups	33
Figure 16: Matching Between Strobe and Data at Receiver for AGP 4X	34
Figure 17: Differential Input Buffer and Required Transfer Characteristic	35
Figure 18: MultiStage Input Receiver	35
Figure 19: On-Chip Signal Metal Layer Cross Section.....	36
Figure 20: Ledge Caused by Imperfect Termination.....	36
Figure 21: System Simulation	37
Figure 22: Strobe Routing to Termination Resistors.....	38
Figure 23: Vref Resistor Divider for AGP 2X Mode	40
Figure 24: Vref Resistor Divider for AGP 4X Mode	41
Figure 25: Noise Coupling Circuit	42
Figure 26: Vref Resistor Divider for AGP3.0 Mode	42
Figure 27: Vref Resistor Divider for All 1.5V Modes	43
Figure 28: Vddq System Power Delivery Model for Graphics Controller Driving	44
Figure 29: Example of a Linear Regulator Circuit.....	45
Figure 30: Example of a Switching Regulator Circuit	45
Figure 31: Worst Case Current Load Waveform (From AGP 4X) To Be Used In VDDQ Power Delivery Simulations.....	49
Figure 32: AGP PCB Layout for ATX Form Factor.....	51
Figure 33: AGP PCB Layout for LPX Form Factor.....	51
Figure 34: AGP PCB Layout For Micro-ATX Form Factor.....	52
Figure 35: Power Plane Etch and Voids	53
Figure 36: Common Clock Timing Budget Between the Graphics Chip, AGP, and Core Chipset	54
Figure 37: Trace Length Routing Ratio	55
Figure 38: Required Trace Length to Minimize Clock Skew.....	57
Figure 39: Interleaved Ground Bondwires for Signal Return and Shielding	60
Figure 40: Cross Section of Bondwires With Signal and Shield Ground Wires In Same Plane.....	60
Figure 41: BGA Package Ballout Examples for AGP2X and AGP4X	61
Figure 42: BGA Package Ballout Examples for AGP8X	61
Figure 43: Example of Skews	63
Figure 44: AGP with One Connector Topology.....	66
Figure 45: Flight Time Measurement.....	70
Figure 46: Signal Quality Measurements	72
Figure 47: Sensitivity Analysis Chart	73
Figure 48: Example of Skew Differences Between 2X Mode and 4X Mode	77
Figure 49: Example of ISI	80
Figure 50: AGP Source Synchronous Topology.....	82
Figure 51: Solution Space Plot	83
Figure 52: Reference Load.....	87

Figure 53: Signal Quality Measurements	88
Figure 54: Ringback Results	89
Figure 55: Minimum Flight Time Results	90
Figure 56: Maximum Flight Time Results	91
Figure 57: Setup Skew Results with 1:4 Trace Spacing	94
Figure 58: Hold Skew Results with 1:4 Trace Spacing	95
Figure 59: Setup Skew Results with 1:3 Trace Spacing	96
Figure 60: Hold Skew Results with 1:3 Trace Spacing	97
Figure 61: Recommended PCB Stackup	98
Figure 62: Upper and Lower Contact Measurement Points	98
Figure 63: Example of Ledge at the Receiver Pin	100
Figure 64: Example of Skew Differences Between 2X, 4X and 8X Mode	101
Figure 65: Microstrip and Stripline Trace	102
Figure 66: Upper and Lower Contact Measurement Points Near Edge Fingers	106
Figure 67: Two Topologies For Interconnect Simulations.....	107
Figure 68: Interconnect Timing Skew Methodology.....	108
Figure 69: Asymmetric Vref and Receiver Timing Windows	109
Figure 70: Data Pattern For Simulations	112
Figure 71: Data and Strobe crosstalk 5-line model example	112
Figure 72: Trace Breakout Example	113
Figure 73: Data and Strobe crosstalk 5-line model example for breakout region	113
Figure 74: Shifting Data Edges	115
Figure 75: Signal Quality Spec.....	116
Figure 76: t_{DRIVE} test waveforms	117
Figure 77: Input drive level test waveform.....	118
Figure 78: Input delay versus input drive level above Vref	118
Figure 79: Input drive level test waveform.....	119
Figure 80: Input delay versus V_{RB} drive time	119
Figure 81: Worst Case Overshoot Results for Microstrip and Stripline Topologies	120
Figure 82: Worst Case Undershoot Results for Microstrip and Stripline Topologies	120
Figure 83: Ringback Results for Microstrip Topology	121
Figure 84: Ringback Results for Stripline Topology.....	121
Figure 85: Timing Skew (Setup) For Microstrip and Stripline Topologies.....	122
Figure 86: Timing Skew (Hold) For Microstrip and Stripline Topologies.....	122
Figure 87: Recommended 4-Layer PCB Stackup	123
Figure 88: Recommended 6-Layer PCB Stackup	123
Figure 89: Measurement point definition	125
Figure 90: Example of ledge at the receiver pin	126
Figure 91: Definition of Edge Placement Accuracy.....	128
Figure 92: EPA Effect On Testing to Specifications.....	128
Figure 93: Moving Tester Strobe to Reduce Escapes.....	129
Figure 94: Settings for Random Measurement Error.....	130
Figure 95: Uncertainty When Using Worst-case Timing	130
Figure 96: Example of Source Synchronous Data Timings	131
Figure 97: Pad to Pin Translations, Driver	133
Figure 98: Pad to Pin Translations, Receiver.....	134
Figure 99: Conventional Tester Timing Calibrations.....	135
Figure 100: Tester Timing Calibrations for Pad Testing	135
Figure 101: Loopback DFT for Source Synchronous Interface.....	136
Figure 102: Eye Diagram Example with Strobe.....	139
Figure 103: Cycle -to-Cycle T_{SU} , T_H Measurement	140
Figure 104: Sample S-parameter Plot for PDN	141
Figure 105: 66 MHz Clock Jitter and Skew Measurement Example	142
Figure 106: 66 MHz Clock Timing at Receiver.....	142
Figure 107: AGP8X Source Synchronous Data/strobe at Receiver	143

Figure 108: AGP8X Logic Analyzer Interface.....	144
Figure 109: AGP8X Logic Analyzer Interface Connector Breakout.....	145
Figure 110: LAI Keep-out Zone.....	146
Figure 111: Probing Baseplate Housing Mounted Around Connector.....	147
Figure 112: Baseplate Housing, Interposer and Anisotropic Material.	148
Figure 113: Interposer board.	149
Figure 114: AGP Add-In Card and Graphics Controller Power Trend.....	150
Figure 115: Airflow Over Add-In Card Components	151
Figure 116: Pulling and Pushing Fan Airflow	152
Figure 117: Adequate Venting Air	153
Figure 118: Inadequate Venting.....	153
Figure 119: Example NLX Air Flow	155
Figure 120: Example NLX Side Panel Vent.....	156
Figure 121: Example NLX Rear Vents.....	156
Figure 122: Example ATX Air Flow	157
Figure 123: Simple Two-Conductor System Used to Explain the Parasitic Matrixes	160
Figure 124: Schematic Used for Calculating Odd and Even Mode Flight Time Differences	162
Figure 125: Schematic Used for Representing a Conductor Pair with a Single Line	162
Figure 126: Three Conductor Even Mode System.....	163
Figure 127: Three Conductor System with Odd Mode and Even Mode.....	164
Figure 128: Five Conductor System.....	165
Figure 129: AGP-8X IBIS Model Vs V-I Spec for Pullup and Pulldown	172
Figure 130: Least Linear with a 1.4 V/ns Slew Rate	173
Figure 131: Least Linear with a 2.8 V/ns Slew Rate	174
Figure 132: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate	174
Figure 133: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate	175
Figure 134: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate	175
Figure 135: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate	176
Figure 136: Least Linear with a 1.4 V/ns Slew Rate.....	176
Figure 137: Least Linear with a 2.8 V/ns Slew Rate.....	177
Figure 138: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate	177
Figure 139: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate	178
Figure 140: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate	178
Figure 141: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate	179

Tables

Table 1: AGP transfer mode definition for Design Guide	15
Table 2: AGP Timing Elements	18
Table 3: Example AGP 1X Mode Delay Components	18
Table 4: Example AGP 2X Mode Skew Components	19
Table 5: Example AGP 4X Mode Skew Components	20
Table 6: Example AGP 8X Mode Skew Components	21
Table 7: Six-Layer Motherboard Stackup Information	33
Table 8: Four-Layer Motherboard Stackup Information.....	33
Table 9: AGP 2X Mode Resistance Values	40
Table 10: Example HF Capacitors Placed Near the Connector on the AGP Card	48
Table 11: Example HF Capacitors Near the Connector on the Motherboard	48
Table 12: Example Bulk Capacitor Stuffing Locations Near the Connector on the Motherboard.....	48
Table 13: Example HF Capacitors Near the MCH on the Motherboard.....	48
Table 14: Example HF Card Capacitors Placed Near the Graphics Controller	49
Table 15: Strobe and Data Groupings for AGP4X and AGP8X	50
Table 16: AGP transfer mode definition for Design Guide.....	62
Table 17: Data and Associated Strobes	64
Table 18: Add-in Card Recommendations.....	64
Table 19: Motherboard Data / Strobe Line Length Recommendations.....	65
Table 20: Control Signal Line Length Recommendations	65
Table 21: Interconnect Impedance and Velocities	68
Table 22: Summary of Simulations	73
Table 23: Minimum Flight Time	74
Table 24: Maximum Flight Time.....	74
Table 25: Summary of Source Synchronous Comparisons.....	75
Table 26: Skew From the Motherboard with 1:1 Trace Spacing.....	75
Table 27: Skew from the Motherboard with 1:2 Trace Spacing.....	76
Table 28: 4X Mode Data and Associated Strobe / Strobe#.....	78
Table 29: Motherboard Recommendations	78
Table 30: Add-in Card Rules	79
Table 31: Control Signal Line Length Recommendations.....	79
Table 32: Board Parameters	85
Table 33: Crosstalk Parameters for 5 mil Data Lines.....	86
Table 34: Crosstalk Parameters for Strobe Lines.....	86
Table 35: Example of the 1024 Combinations of Interconnect Parameters	91
Table 36: Reduced Set of Interconnect Parameters to be Simulated	92
Table 37: Example of Simulation Combinations.....	93
Table 38: 8X Mode Data and Associated 1 st and 2 nd Strobes	102
Table 39: 8X Mode Motherboard Rules	103
Table 40: 8X Mode Add-in Card Rules	104
Table 41: 8X Mode Control Signal Line Length Recommendations	105
Table 42: Package Trace Matching Example	106
Table 43: 8X Mode Timing Budget	109
Table 44: AGP Connector Simulation Model Pin Map	111
Table 45: Simulation Sweep Matrix.....	114
Table 46: Skew Calculation Criteria	114
Table 47: Recommended Time Domain Test Equipment.....	137
Table 48: Results of Single Line and Coupled Line Simulations Showing Agreement Between Two-Conductor and One-Conductor Flight Times for Both Even and Odd Modes	162
Table 49: Results of Multi-line and Single Line Flight Time Simulations Showing Agreement Between 3-Conductor and Equivalent Single Conductor Models	166
Table 50: AGP-4X Buffer I-V Curves	167

Table 51: AGP-4X Diode Clamp I-V Curves	168
Table 52: AGP 8X Buffer V-I Curves	169

Preface

Document Structure and Outline

This design guide discusses design examples for Accelerated Graphics Port (AGP) 1X, 2X, 4X, and 8X modes which also include AGP low voltage designs and 0.8 volt swing designs. The *Accelerated Graphics Port Platform Design Guide*, Revision 1.0, published in August 1998, has been incorporated into this document.

The topics covered in this design guide are not specifications. For specifications, refer to the *Accelerated Graphics Port Interface Specification, Revision 2.0* and *AGP3.0 Interface Specification, Revision 1.0* available on the AGP Implementer's Forum Web site (www.agpforum.org) or on the Intel AGP Web site (<http://developer.intel.com/technology/agp>).

Content

The initial release of this document centered around electromechanical implementation issues. The first update added thermal design concerns. Revision 1.0 incorporated the prior document and added the AGP 4X guidelines. Revision 1.5 incorporates all of Revision 1.0, includes clarifications and new information for AGP 4X, and adds AGP 8X guidelines. This revision updates the design guidelines and simulation data. Refer to the Intel AGP Web site to view all Engineering Change Notices (ECNs) for the current revision.

Feedback

We welcome comments, feedback, and suggestions. Please submit them on the AGP Implementer's Forum web site (www.agpforum.org).

1 AGP Electrical Design Considerations

1.1 AGP Overview

The AGP physical interface is a point-to-point topology using 1.5 volt or 3.3 volt signaling. The baseline performance level for AGP uses a 66 MHz clock to provide a peak bandwidth of 266 MB/s.

A double-clocking data technique is used to achieve twice the baseline bandwidth. Thus, AGP 2X mode provides a peak bandwidth of 533 MB/s. AGP 2X mode is a superset of the 1X mode.

The AGP 4X mode provides high performance levels with a peak bandwidth of 1066 MB/s. This bandwidth is achieved by using a quad clocked data transfer methodology, which allows four times as much data to be transferred every 66 MHz clock cycle. With the demand for increasing data transfer rates in high-speed systems, the process technologies must continually improve. AGP 4X requires a process technology, which is 0.35 μm or better in order to meet the tighter timing specifications. AGP 4X mode is a superset to the 1X and 2X modes; thus, all components supporting AGP 4X must also support 1X and 2X modes.

The AGP 8X mode doubles the data transfer rate yet again to 2.1 GB/s peak bandwidth. The data is octal-pumped from the common clock frequency of 66 MHz. In order to achieve these higher data rates across the standard AGP connector, AGP 8X mode uses a parallel-terminated bus with low signal swings of about 0.8 volts, ground referenced. The AGP 4X mode 1.5 volt I/O power supply is used for backwards compatibility for designs which wish to be universal low voltage. It is likely that a 0.25 μm or better technology will be required to meet the tight timing specifications.

In this chapter, the terms 1X, 2X, 4X, and 8X modes refer to a specific bus electrical signaling and transfer speed according to the following table:

Table 1: AGP transfer mode definition for Design Guide

Mode	Transfer speed	Electrical signaling	Definition reference
1X mode	66MT/s	3.3V or 1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
2X mode	133MT/s	3.3V or 1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
4X mode	266MT/s	1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
8X mode	533MT/s	0.8V signal swing.	<i>AGP3.0 Interface Specification, Revision 1.0</i>




Figure 1: System Overview

A typical system physical configuration is shown in Figure 1 where the system core logic is in a chipset on the motherboard. The graphics controller is on a daughter card and attaches to the AGP interconnect via an AGP edge connector. Other arrangements are possible (graphics controller “down” on the motherboard), but the AGP Specification or this design guide does not cover those. The AGP electrical design portion of this document covers three main components: transmitter, interconnect, and receiver as shown in Figure 2. It focuses on electrical issues when designing these components.




Figure 2: AGP Design Components

1.1.1 System Timing Overview

The AGP timings can be segmented into various time domains as shown in Figure 3. The outer loop of the transmitter and receiver both operate from a common AGP clock, and the outer loop controls are specified relative to this clock as in the 1X mode. The inner timing layer, referred to as the inner loop, specifies timing relationships for the reliable transfer of data from the transmitter's output latches to the receiver's input latches as in the 2X, 4X and 8X transfer modes. The inner loop timings use additional source synchronous timing signals to realize the data transfer. Refer to the *Accelerated Graphics Port Interface Specification, Revision 2.0* and *AGP3.0 Interface Specification, Revision 1.0* for more details.

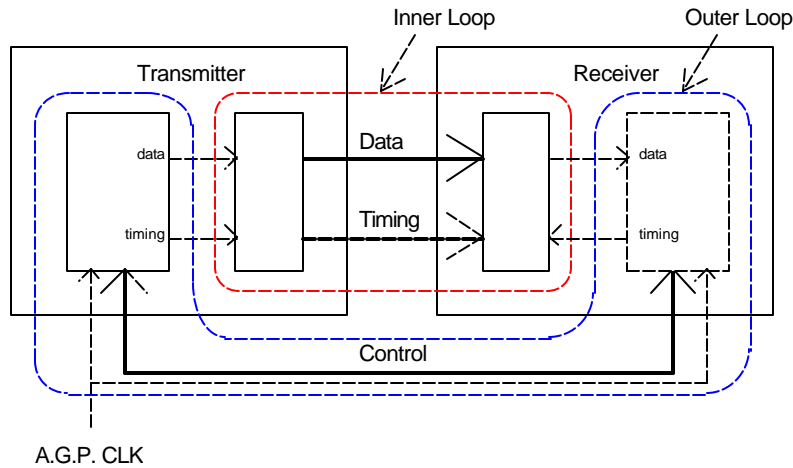


Figure 3: AGP Time Domains

The timings for AGP 1X mode are based on a common clock and are driven by I/O delays and the clock skew between the AGP components. The delays of AGP 1X add up to one AGP clock period (15 ns), the maximum time allowed for a data transfer. The AGP 2X mode, 4X mode and 8X mode timings are based on the summation in skew between the strobe and data lines for the transmitter, interconnect, and the receiver. The skews of AGP 2X add up to a quarter of an AGP clock period (3.75 ns), the nominal phase difference between a data transition, and a strobe transition. Similarly, the skews for AGP 4X mode add up to 1.875 ns and the skews for AGP 8X mode add up to 937.5 ps. A summary of the key elements is given in Table 2. Table 3, Table 4, Table 5, and Table 6 show a detailed breakup of the skew components of the transmitter, interconnect, and receiver for AGP 1X, 2X, 4X, and 8X modes. These timing budgets and the factors influencing them are discussed in the following sections.

Table 2: AGP Timing Elements

	AGP 1X (Delay)	AGP 2X (Skew)	AGP 4X (Skew)	AGP 8X (Skew)	Units
Clock Skew	1000	NA	NA	NA	ps

	AGP 1X (Delay)	AGP 2X (Skew)		AGP 4X (Skew)		AGP 8X(Skew)		Units
		Setup ²	Hold ²	Setup ²	Hold ²	Setup ²	Hold ²	
Transmitter	6000 ¹	2050 ³	1850 ³	925 ³	725 ³	410 ³	460 ³	ps
Interconnect	2500	700	900	550	450	442.5	267.5	ps
Receiver	5500	1000	1000	400	700	85	210	ps
Total	15000²	3750	3750	1875	1875	937.5	937.5	ps

Note 1: Data timings

Note 2: Includes clock skew. Clock skew is not an element of AGP 2X, 4X, or 8X timings.

Note 3: The transmitter skew is from the internal clock. Setup skew = 0.5bit time – tDVb and hold skew = 0.5bit time – tDVa.

Table 3: Example AGP 1X Mode Delay Components

Element	Delay Component	AGP 1X	Units
Transmitter	Internal clock delay	1000 ²	ps
	Clock jitter	250	ps
	Clock to output delay	4000	ps
	SSO pushout	500	ps
	Design / Tester Guardband ¹	250	ps
	Total Delay at Transmitter	6000	ps
Interconnect³	Delay at Interconnect	2500	ps
Receiver	Input buffer delay	1500	ps
	Input logic and routing delay	4250	ps
	Data register setup time	500	ps
	Design / Tester Guardband ¹	250	ps
	Internal Clock Delay	-1000 ²	ps
	Total Delay at Receiver	5500	ps

Note 1: This example margin is to cover testing or design uncertainty. See sections 3.1.1 and 3.1.2 for more information on tester and design guardbands.

Note 2: The internal clock delay acts to delay outputs, but assists on input setup times.

Note 3: Interconnect delays and skews are for motherboard and add-in card combined.

Table 4: Example AGP 2X Mode Skew Components

Element	Skew Component	Setup	Hold	Units
Transmitter	Data / strobe clock jitter	250	250	ps
	Data / strobe clock duty cycle	150	150	ps
	Internal data clock skew	100	100	ps
	Buffer delay matching	150	150	ps
	Rise / fall time matching	750	750	ps
	SSO pushout	500	200	ps
	Design / Tester Guardband ¹	150	150	ps
	Total Skew at Transmitter	2050	1850	ps
Interconnect ²	Skew at Interconnect	700	900	ps
Receiver	Strobe to data path skew	650	650	ps
	Strobe routing skew	200	200	ps
	Design / Tester Guardband ¹	150	150	ps
	Total Skew at Receiver	1000	1000	ps

Note 1: This example margin is to cover testing or design uncertainty. See sections 3.1.1 and 3.1.2 for more information on tester and design guardbands

Note 2: Interconnect delays and skews are for motherboard and add-in card combined.

Table 5: Example AGP 4X Mode Skew Components

Element	Skew Component	Setup	Hold	Units
Transmitter	Clock Uncertainty			
	Internal clock delay	NA	NA	ps
	Clock to output delay	NA	NA	ps
	Data / strobe clock jitter	150	150	ps
	Data / strobe clock duty cycle	75	75	ps
	Internal data clock skew	25	25	ps
	Buffer delay matching	25	25	ps
	Rise / Fall time matching	150	150	ps
	SSO data pushout	350	0	ps
	SSO strobe pushout	0	150	ps
	Design / Tester Guardband ¹	150	150	ps
	Total Skew at Transmitter	925	725	ps
Interconnect ²	Skew at Interconnect	550	450	ps
Receiver	Strobe to data path skew	100	100	ps
	Strobe routing skew	50	50	ps
	Setup and Hold	100	400	ps
	Design / Tester Guardband ¹	150	150	ps
	Total Skew at Receiver	400	700	ps

Note 1: This example margin is to cover testing or design uncertainty. See sections 3.1.1 and 3.1.2 for more information on tester and design guardbands

Note 2: Interconnect delays and skews are for motherboard and add-in card combined.

Table 6: Example AGP 8X Mode Skew Components

Element	Skew Component	Setup	Hold	Units
Transmitter clock Transmitter driver	Jitter	150	150	ps
	Data / strobe clock duty cycle	50	50	ps
	Distribution variance in the I/O ring	10	10	ps
	Tco mismatch between buffers	25	25	ps
	Rise/fall time variance	50	100	ps
	SSO pushout	75	75	ps
	Design guard band ¹	50	50	ps
	Total Skew at Transmitter	410	460	ps
Interconnect ²	Total Interconnect Skew	442.5	267.5	ps
Receiver	Setup or hold requirements	15	140	ps
	Si mismatch	20	20	ps
	Design guard band ¹	50	50	ps
	Total Skew at Receiver	85	210	ps

Note 1: This example margin is to cover design uncertainty. See sections 3.1.1 and 3.1.2 for more information on tester and design guardbands.

Note 2: Interconnect delays and skews are for motherboard and add-in card combined.

1.2 Transmitter

Most of the components of the AGP 1X timings should be familiar to designers with experience in I/O design, especially PCI I/O design. The AGP 1X timings will not be discussed in detail here. However, it is important to be sure that all causes of delay and skew are included, such as clock jitter and simultaneous switching outputs (SSO) pushout delay.

All the components of the AGP 2X, 4X and 8X timings need to be carefully considered since there is less timing margin in these cases, and several small contributions add up fast. The magnitudes of the component timings are smaller since they are generally the mismatched delays (skew) of two paths. Since the paths are made of the same or very similar circuits, the skew can be minimized if care is taken. The skew components will be discussed in detail along with the factors influencing them and how to minimize skew.

1.2.1 Internal Clock Generation and Distribution

For all AGP modes 2X and above, the transmitter is assumed to use a PLL to produce a 50% duty cycle clock. For AGP 2X, the PLL generates a 133 MHz clock to produce the required relative timings of data and strobe. For AGP 4X mode, a 266 MHz clock is needed to capture twice as many data packets within one cycle as compared to the AGP 2X mode. For AGP 8X, a 533 MHz clock is required. The PLL is the primary source for the data/strobe clock jitter and clock duty cycle skew components. Duty cycle degradation can occur as the clock travels through the interconnect and receiver. A 50% duty cycle clock provides the best balance of data to strobe setup and hold times. Any error in the clock duty cycle or any clock jitter reduces the time from the data to strobe (setup time) or strobe to data (hold time).

The PLL has an intrinsic jitter caused by its clock source, supply noise, crosstalk with digital signals on the same chip, and sometimes from sources within the PLL itself. Jitter can be reduced by providing a low jitter clock source and by proper isolation and bypassing of the power to the PLL. The PLL may

output a good clock duty cycle with the voltage controlled oscillator (VCO) running at 133 MHz for AGP 2X mode, 266 MHz for AGP 4X mode, or 533MHz for AGP 8x. If the duty cycle is not close to 50%, it may be necessary to run the VCO at twice the above-mentioned frequencies and then divide it by two. Clock duty cycle is also affected by internal loading and buffering of the clock.

A balanced clock distribution is recommended. The strobe signals and clock signals must be routed with their associated data group. For example, in AGP 4X, a Strobe/Strobe# pair is provided for every 16 AD bus signals. These should be placed in the center of their respective data groups to minimize data strobe timing skew. The clock signals should be routed out from this same point to also reduce clock skew. Avoid auto-routing the clock signals to prevent exceeding the skew budget in the interface. The clock path should be designed and validated against a strict timing budget like those shown in Table 4, Table 5, and Table 6 to ensure signal fidelity.

The 66 MHz or 133 MHz clock can be used as the core data clock. The data transmit clock can use the 66 MHz clock for AGP 1X mode, 133 MHz clock for AGP 2X mode, 266 MHz for AGP 4X mode, or 533 MHz for AGP8X.

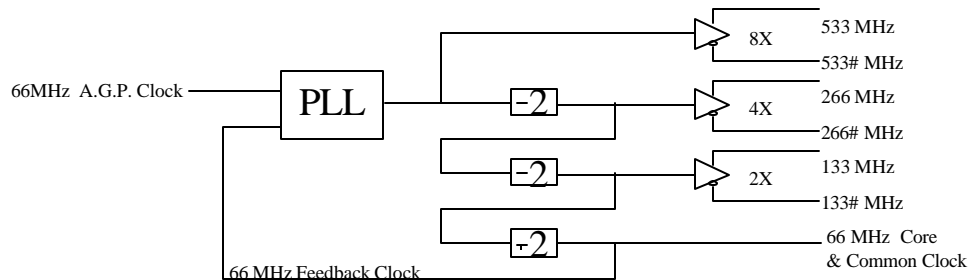


Figure 4: Clock Design

When designing the clocking scheme, the main objective is to ensure matching between the clock signals. The data transmit clocks and the strobe clocks differ in phase; however, identical buffers should be used for data signals and strobe signals to reduce mismatch. A balanced routing topology will aid in matching the trace lengths within the groups to minimize skew.

For AGP, two methods for clocking to the data and strobe buffers are shown in Figure 5. One implementation is to use both the true and complement 133/266/533 MHz clocks for data and strobe on the interface. The benefit to this method is it allows strobe and data to use identical buffers timed off of the rising edge. The drawback with this approach is the loading on each clock is different which in turn creates duty cycle variations due to the delay variation. Timing adjustment is critical and must be taken into account. The mismatch in load can be balanced by dummy loads in the interface cells that do not use that particular phase of the clock. Identical clock drivers and routing will null out most of the skew.

Another implementation is to use a single clock source with the data transfers clocked on the rising edge and the strobes clocked on the falling edge. The benefit is using only a single clock; however, it is essential to match the data's rising edge flip-flop delay to the strobe's falling edge flip-flop delay. Again, a 50% duty cycle clock is critical.

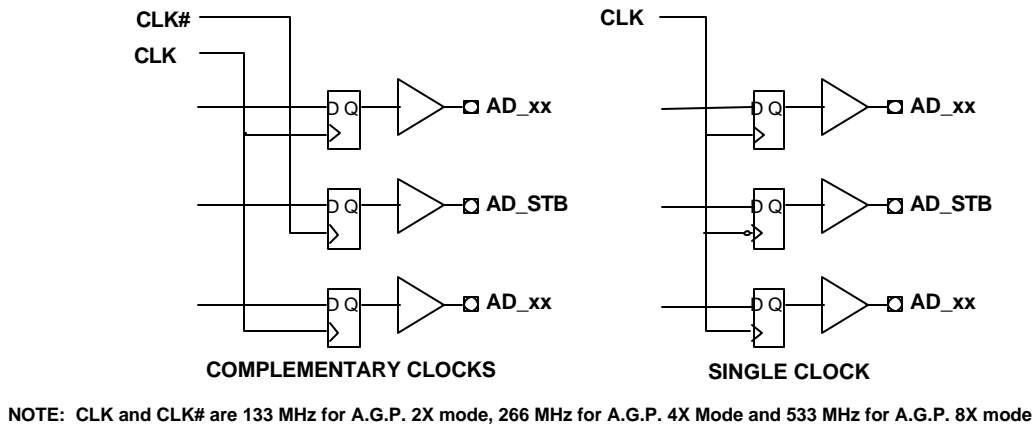


Figure 5: Clocking Methodologies

When designing a system that includes AGP 2X, 4X and 8X, it may be easier to use a 533 MHz clock (and its inverse) for all AGP modes, depending on the one in use, rather than switching clock speeds. State machines running at 533 MHz can be designed to implement AGP 2X and 4X serializing and de-serializing of data, as well as generating strobe signals. In AGP 8X, mode 533 MHz can be used to generate data, while 533# MHz can be used to generate strobe. (See Figure 6.)




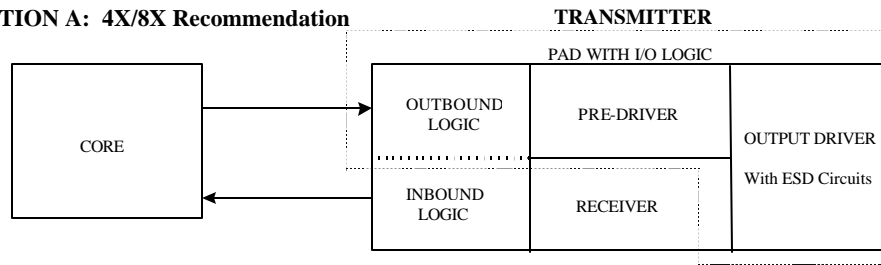
Figure 6: Clock Design for AGP 8X/4X/2X System

1.2.2 Output Buffer Considerations

1.2.2.1 Buffer Architecture

Placement of the buffer components is critical in order to control the skew within a group of data signals with respect to their strobes. The I/O buffer consists of inbound (from I/O to the core) and outbound (from core to I/O) latches, multiplex, and flip-flops. When designing the pad ring, there are two options when integrating the I/O logic. The first option is to design the I/O logic and registers, pre-driver, and the output driver into the pad ring, as shown in Figure 7, Option A. This method is recommended for the AGP 4X and 8X mode. The clock to output time and the variations between data and strobe are minimized with this implementation. Another option (Figure 7, Option B) is to custom route the I/O logic between the core and the pad. The logic should be placed as close as possible to the pad to reduce mismatches/skew due to RC delays. There are additional clock to output delay and variations incurred due to the distance between the logic and pad when implementing option B. This option also requires matching signals between the I/O and logic pads. This method can be used for the AGP 2X mode; however, the designer should be aware of the increased driver skew created by this type of implementation.

OPTION A: 4X/8X Recommendation



OPTION B: 2X Recommendation




Figure 7: Buffer Configurations

Figure 8 shows greater details of the data and strobe buffer architecture. The same buffer design is useful for data and strobe signals.

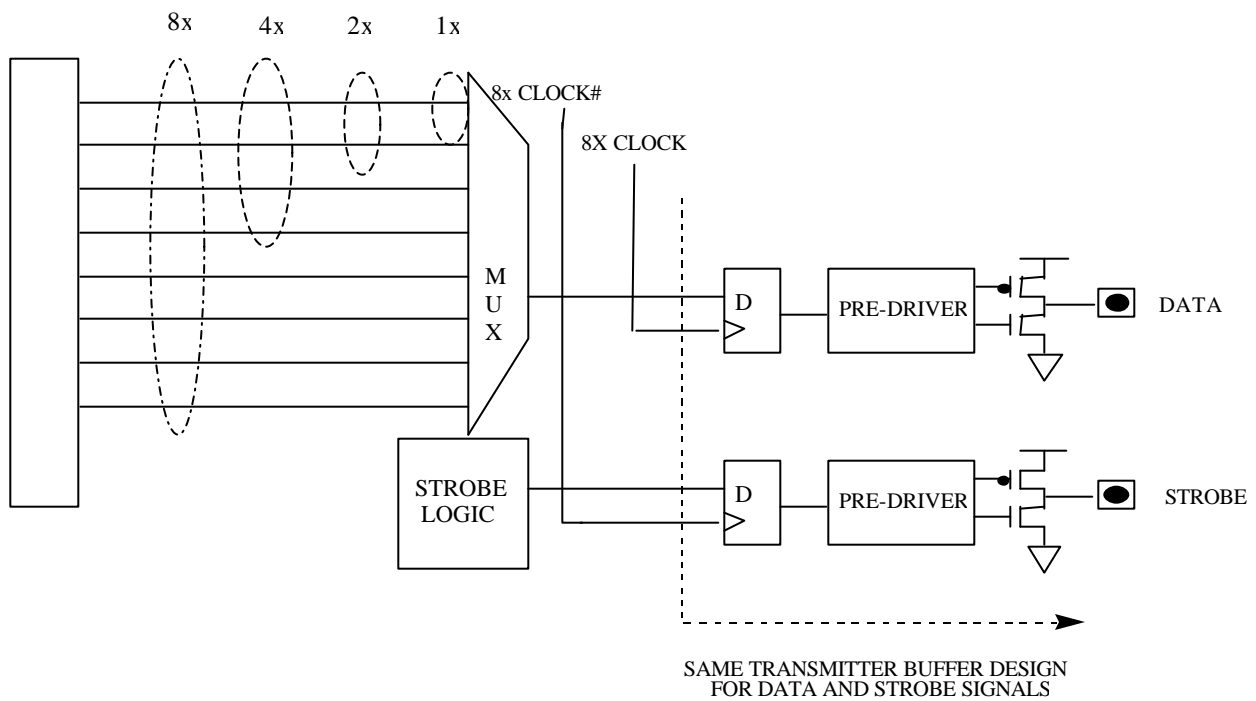



Figure 8: Data and Strobe Buffers

1.2.2.2 Buffer Design

Typical CMOS techniques can be used for the AGP 2X mode buffer design. However, these techniques are not sufficient to meet the timing requirements of AGP 4X or AGP 8X, which is the focus of this section.

The buffer itself needs to be carefully designed to balance the delay mismatches of the rising and falling edges. Any delay mismatch reduces data setup and hold time margin. One way to reduce rise and fall time variations is to use non-minimal channel length pre-drivers with impedance compensated output stages, thus reducing power, voltage, and temperature variations. Process, voltage, and temperature variations can add up to $\pm 50\%$ variability in parameters such as the output impedance, R_{out} , of the output buffer. Output impedance compensation can reduce the variability from $\pm 50\%$ down to $\pm 20\%$. $R_{out} \pm 20\%$ covers all effects including simultaneous switching outputs (SSO). The recommendation is to design to $R_{out} \pm 10\%$. This ensures that buffer impedance variation stays within $\pm 20\%$ when taking SSO, power supply fluctuations, etc., into account.



1.2.2.2.1 Digital Compensation

For AGP 4X and AGP 8X mode, digitally compensated buffers are recommended. Figure 9 shows an example of a digital compensation scheme. The reference generator uses a resistor bridge to select transistor leg sizes, by comparison, against a reference resistor value which is proportional to the target buffer impedance. The reference resistor, R_{Z_0} , can be an external or internal resistor depending upon tolerance requirements. The legs are enabled until division between the reference resistor, R_{Z_0} , and the enabled legs equals the voltage at which the buffer impedance is evaluated. For AGP 4X, this voltage is V_{ref} , the same voltage as that described in Section 1.4.4.2. For AGP 8X, the voltage is the target swing voltage V_{swing} of 0.8 volts nominal. Only the circuit for the pull-down legs of the driver is shown, but analogous methods can be used to set the buffer strength of the pull-up devices.

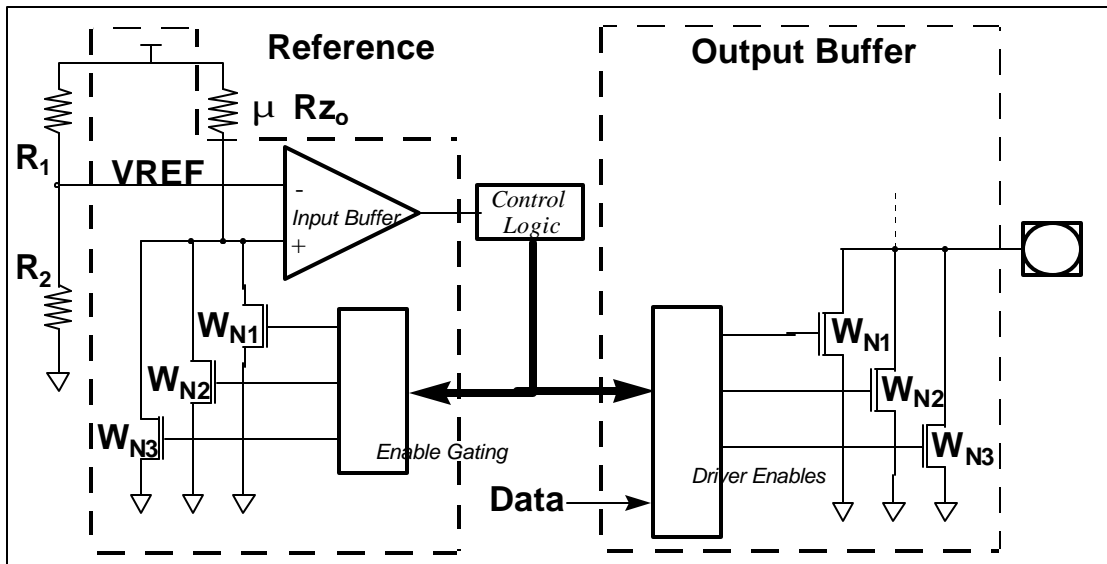


Figure 9: Digital Impedance Compensation

Buffer impedance variation can be compensated by using 3-bit or 4-bit binary weighted digital compensation techniques with one leg always on. Each leg is binary weighted, allowing for a simpler design, uniform coverage with fewer steps, and greater efficiency.

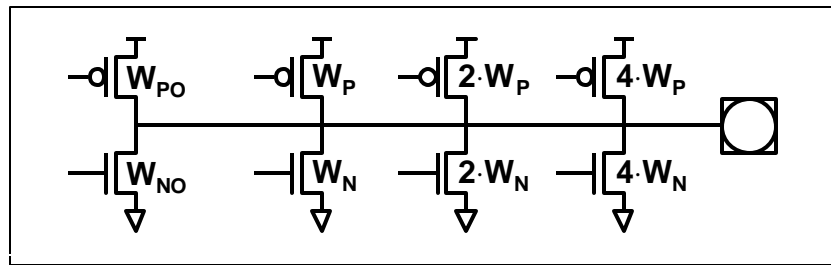


Figure 10: Binary Encoding

One benefit from using digital compensation is that compensation information can be easily distributed across the pad ring as a DC signal. By designing the output buffer impedance to tighter constraints in a 5-bit binary weighted scheme, for example, full range compensation can be achieved with 3% variation but with the added cost of validating compensation functionality.

The digital compensation scheme should be carefully designed and evaluated for output impedance error due to quantization and drift of temperature and V_{ddq} (V_{cc}). In order to avoid such drift, the compensation information to the output buffer may require periodic updates to maintain required tracking of temperature and voltage. Updates are susceptible to glitches, for example, when the compensation information transitions from 1000 to 0111. The methodology to update impedance values needs to be carefully managed. Preferably, the output buffer should be updated at boot up time and once every few hundred microseconds when the interface is idle or not driving the bus. AGP 8X has special cycles for update cycles that are described in the *AGP3.0 Interface Specification, Revision 1.0*.

1.2.2.2.2 Low Curvature Buffers

A buffer with more linear V-I characteristics improves signal quality by absorbing reflected signal waves and reducing ringback. Linear buffers allow all reflections to see the same “reflection coefficient” when the wave returns to the driver. High curvature (a saturated V-I characteristic) means the reflection coefficient is voltage-dependent; therefore, waves at slightly different voltages will generate vastly different reflections, contributing to uncertainties in flight time and skew. This uncertainty is greatly reduced with more linear buffers.

Adding resistors in the output path makes the buffer more linear. A series resistor creates a voltage drop that helps keep the transistor out of the saturation region at larger pad voltages. Series resistance, as shown in Figure 11, largely determines the buffer impedance.

There is a drawback to this implementation. The series resistance weakens the buffer and a larger buffer is needed to obtain a given drive strength and equivalent impedance. Therefore, buffer characteristics chosen for AGP 4X and 8X are a balance of improved signal quality and drive strength loss.




Figure 11: Increasingly Linear Buffer Curves

1.2.2.2.3 Measuring Curvature of I/V Curves – AGP 4X

The minimum and maximum drive characteristics of the AGP 4X output buffers are defined by V/I curves. These curves are the DC buffer characteristics. They are measured from the device pin by sweeping the pad voltage from ground to Vddq. (A sweep of the V/I characteristics outside of the ground and Vddq rails can also be done to evaluate the buffer clamp characteristics, which are also covered by the specification.) The required curves are more linear than typical CMOS curves with a less pronounced saturation region “knee.” The 48 Ω and 73 Ω load lines shown in Figure 12 have three purposes in measuring the buffer characteristics and curvature:

- First, the 48 Ω load line represents the lowest recommended effective interconnect impedance. This represents the drive level necessary to be reached within the output delay time t_{VAL} in order to get a good drive level at the receiver. The partitioning of the t_{VAL} between clock distribution, logic, and output buffer parameters is not a matter for this specification.
- Second, the instantaneous impedance of the V/I curve at the point where the 48 Ω load line intersects the curve defines the limits of instantaneous impedance of the curves.
- Third, the ratio of the instantaneous impedances, where any curve within the range crosses the 48 Ω line to the instantaneous impedance where it crosses the 73 Ω line, defines the curvature of the V/I curve. It is important for signal integrity (especially skew and ringback) that the buffers lie within the range specified.

To measure the curvature of the V/I curve within the 48-73 Ω operating impedance range (see the shaded area in Figure 12), first determine the instantaneous impedance by measuring the slope at the 48 Ω load line intercept (point 1 in Figure 12). This is then repeated for the 73 Ω load line (point 2 in Figure 12). The curvature is then determined by the ratio of these two instantaneous impedances. The slope of the curve can be calculated from simulation or from measurement data using a linear approximation of delta voltage over delta current. A maximum of 10 mV steps is recommended when calculating the slope.

Adherence to these curves should be evaluated at worst-case system conditions, including the tolerances on any components used to target the buffer impedance compensation (i.e., the pull-up resistor proportional to Z_o in Figure 9). The worst-case conditions should include environmental effects, including external and internal power supply variations due to simultaneous switching outputs (SSO). The minimum curves should be evaluated at minimum Vddq and high temperature. The maximum curve test points should be evaluated at maximum Vddq and low temperature.




Figure 12: AGP 4X Pull Down/Pull Up V/I Curve

1.2.2.2.4 Measuring Curvature of I/V Curves – AGP 8X

When operating in AGP 8X mode, the pull-down driver impedance should be set at 50 ohms. The driver impedance is measured as $V_{\text{swing}} / I_{\text{out}}$ or $0.8v / I_{\text{out}}$, where I_{out} is the output current of the driver at 0.8v. The impedance linearity should be within 10% of the target impedance. The impedance linearity is calculated as $Z_{\text{at origin}} / Z_{\text{buffer}}$, where Z_{buffer} is $V_{\text{out}} / I_{\text{out}}$ of the buffer. Notice, that the important point here is that for estimating transmission line effects, the termination impedance is calculated as $V_{\text{out}} / I_{\text{out}}$ from the I-V curve (and not the instantaneous impedance $dV_{\text{out}} / dI_{\text{out}}$). Thus the linearity measurement calculates the percentage variation of $V_{\text{out}} / I_{\text{out}}$ throughout the entire V_{out} range. The pictures below illustrate the point.




Figure 13: Linearity measurement

1.2.2.2.5 Output Buffer Reference Loads

It is recommended that a transmission line model be used for evaluating the performance of the output buffers, especially for AGP 4X. The model should be designed with 16 bits of data along with Strobe/Strobe# signals. The power and ground pads should be placed between the **AD** pins exactly as they occur on the die. The recommended transmission line characteristics are given in the *AGP Interface Specification, Revision 2.0*. If a transmission line model is not available, a simple interconnect (RLC) model can be substituted to evaluate and design output buffers, and to study the rise and fall time requirements. Regardless of the load modeling, the output buffer model should include the package parasitic capacitances, inductances, and resistances to study the effects of SSO noise and subsequent degradation of rise and fall times. A full description of these techniques is beyond the scope of this

document. Note that a transmission line load is not the standard load into which the output specifications are measured for AGP 8X. A simple 50 Ω resistor is used for that. See the *AGP3.0 Interface Specification, Revision 1.0* for details.

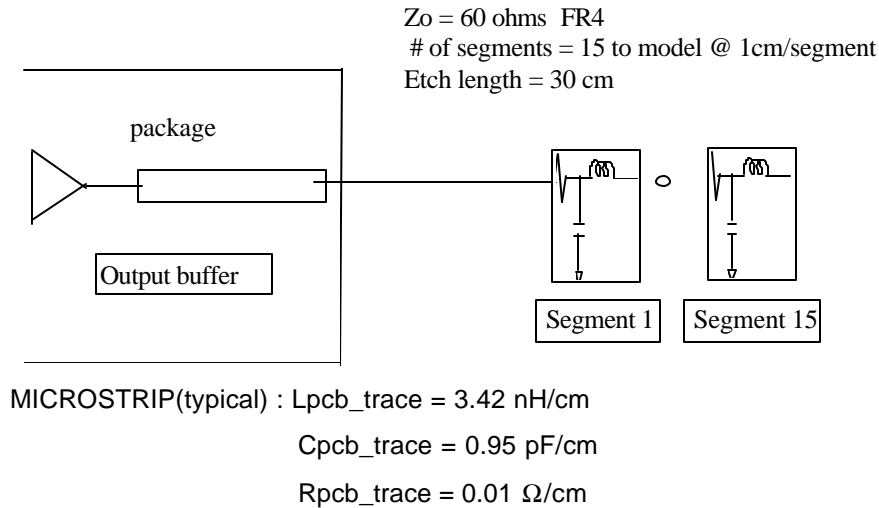


Figure 14: Output Buffer Reference Network

1.2.3 Simultaneous Switching Output (SSO) Noise and Pattern Effects

Simultaneous switching output (SSO) occurs when a group of outputs are switched in the same direction at the same time. This is typically one of the bigger components of transmitter skew. The output delay of the buffer slows a bit due to internal ground or power line voltage drop (resistive and inductive effects). While all other skews are equally likely to reduce setup or hold time, SSO is different in that it mostly causes data delay. Subsequently, SSO has dominant effects on data setup time.

Increasing the number of ground and power connections to the output buffers (reducing the supply resistance and inductance) can reduce SSO pushout. However, at some point, this adds to the device packaging cost. SSO pushout can also be reduced by designing the buffers to the weak end of the buffer strength specification with the largest allowed rise and fall transition time at the cost of increased skew sensitivity, especially for AGP 4X and 8X. Slew rate compensation can also be added to the buffer to limit it to the lower half of the recommended range. These methods reduce the buffer di/dt and, consequently, the inductive component of power supply voltage drop. Adding capacitance from V_{ddq} to ground on the die can also reduce the impact of buffer switching noise. It should be noted that the die capacitance and power supply inductance forms an LC network. If the resonant frequency of this network is near any of the data transmission frequencies, the impedance of the power delivery will significantly increase, causing larger power drops and greater timing pushout. Adding die capacitance will detune the LC network.

All buffer designs will suffer some level of SSO pushout. One such consequence is that data hold time is always longer than data setup time. Because of this asymmetry, the AGP 2X and 4X data hold time after strobe (t_{Dva}) is specified to be 200 ps more than the data setup time before strobe (t_{Dvb}). This establishes 200 ps as the minimum allowance for SSO pushout. For AGP 8X, the impact on the driver is 75ps for setup and hold, due to the SSO pushout.

In AGP 4X, the Strobe/Strobe# signals are 90° out of phase (lagging) with respect to the data. The effects of power and ground noise on the strobe signals can cause strobe pushout. The Strobe/Strobe# signals should be adjacent to one another and shielded with separate power and ground pads on the die. If there is significant internal power bus ground hop or power supply droop during SSO on the data lines, it is recommended that the Strobe/Strobe# power and ground connections be isolated from the data power. This prevents the power supply glitches from disturbing the strobe and possibly causing a glitch. A glitch on the strobe signal may be interpreted as a strobe pulse when it arrives at the receiver. In AGP 8X mode, this is unnecessary, since better termination and DBI (Dynamic Bus Inversion) significantly reduces signal noise. However, a 2-to-1 signal to ground, 4-to-1 signal to power ratio, and heavy power-

to-ground decoupling (about 100pF per data pin) on die must be maintained to reduce power supply noise and SSO related effects.

Data patterns can be used to evaluate signal integrity. The following is an example reference data pattern to capture overall signal integrity and can be used to drive the reference network shown in Figure 14. This specific pattern is used to evaluate inter-signal interference (ISI) effects based on pattern dependencies:

Bit Number	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Pattern	0	1	0	1	0	1	0	1	0	0	1	1	0	0	1	1	0	0	1	1	1	0	1	1	1	0	0	0	1	0	0	0

The following are simulation patterns to study the effects of SSO push out and coupled noise into a victim line:

- All buffers switching low to high or high to low (0000 to FFFF or FFFF to 0000) and observing STB/STB# for any noise coupling (since data and strobes are phase shifted by 90°)
- All buffers switching low to high or high to low (0000 to FFFF or FFFF to 0000)
- All buffers switching low to high except one buffer remains low (0000 to 1110)
- All buffers switching high to low except one buffer remains high (1111 to 0001)
- All buffers switching low to high except one buffer switching high to low (0001 to 1110)
- All buffers switching high to low except one buffer switching low to high (1110 to 0001)

Simulation waveforms at the pad and package should be captured as a reference for comparison to actual waveforms on a board. The noise on the power supplies should also be captured for comparison to actual waveforms. The simulations with all lines switching but one can also be used to monitor internal power line noise.

1.2.4 Transmitter Floorplan

The goal of a good layout is to avoid variations or mismatches that affect the inner loop timings at the transmitter clocks, data and strobe pad, and inbound and outbound logic. The following floorplan recommendations are an example of a good layout that helps to meet AGP timing goals in both the AGP compliant graphics device and the core logic. This kind of layout is especially critical to meet AGP 4X and 8X timings. These are design recommendations and not specifications.

- A ground pin should be placed between every two **AD** pins in order to minimize crosstalk and SSO. There should be at least a power pin per 3-4 **AD** signals nearby to stabilize the power distribution.
- A Strobe/Strobe# pair is provided for every 16 AD bus signals. These should also be placed in the center of their respective data groups to minimize skew.
- Auto-routing the clock signals is not recommended in order to avoid exceeding the skew budget.
- The AGP 2X, 4X, and 8X mode clocks should be routed from the center of the strobe group.
- The Vref signals and clock signals should be shielded with Vss to reduce noise and jitter, respectively.
- The power and ground signals should be laid out based upon the package and multi-row bonding within the package (power and ground "swizzle"), which dictates how adjacent signals decouple against each other.

1.3 Interconnect

In order to meet the tighter timing requirements, the interconnect mismatch must also be minimized. An ideal case is to have the data and strobe signals arrive at the receiver with no change in their timing relationships. In an actual environment, there are trace length mismatches associated with routing multiple signals. For AGP 2X mode, the data lines should be kept to within ± 0.5 inches of their respective strobe on the add-in card and kept at +0.0, -0.5 inches from their respective strobe on the motherboard. In AGP 4X mode, the trace line length mismatch with respect to the strobe lines to any data line in their group should be less than ± 0.25 inch on add-in cards and ± 0.125 inch (7.25 inch lines) on the motherboard. For AGP 8X, the trace line length mismatch with respect to the strobe lines should be ± 25 mils on add-in cards and ± 25 mils on motherboards.

Because the strobe serves many data inputs, its input buffer is more heavily loaded. This usually results in more delay and must be compensated in the data path. Each strobe control spans half the data inputs, and the skew of the strobe arriving at each data input should be well matched. The strobe buffer should be located physically in the center of the data pins it serves. Strobe input buffer loading can be reduced by providing one or more input buffers distributed through the interface, all driven from the same strobe input pad. While this can reduce loading effects and skews, care has to be taken that the capacitive load of the strobe (including packaging) is within the range of ± 1.0 pF for AGP 4X or ± 0.5 pF for AGP 8X as compared to the capacitive load of all data pins connected with that strobe.

Crosstalk is caused by the coupling of a signal to adjacent signals on the board. It can cause speed degradation and glitches to occur in an otherwise stable signal. Crosstalk can be reduced by increasing the distance between signals when laying out the board and shielding the signals with ground planes. The crosstalk can also be reduced by increasing the ratio of the thickness of the prepreg to trace spacing. Figure 15 shows recommended board stack-ups for 4 and 6 layer boards with thin dielectric layers that help reduce crosstalk. Table 7 and Table 8 give the board and routing characteristics for these stack-ups for AGP-4X and AGP-8X systems. Refer to Section 1.6 for additional guidelines.

See Chapter 2 for crosstalk and ISI effects and timing penalties.

AGP Design Guide, Rev. 2.0




Figure 15: Example PCB Layer Stackups

Table 7: Six-Layer Motherboard Stackup Information

	AGP 4X Mode	AGP 8X Mode
Material	FR-4	FR-4
Impedance	60 Ω \pm 15%	60 Ω (microstrip), 56 Ω (stripline) \pm 10%
Board Thickness	62 mils \pm 10%	62 mils \pm 10%
Signal Layer Trace Width	6 mils	6 mils
Component Layer Trace Width	5 mils	5 mils
Breakout	5/5 spacing for BGA breakout and for routing in and out of connector	5/5 spacing from BGA breakout (500 mils) and for routing in and out of connector (200 mils)
Spacing (Data-to-Data)	5/15 mils	5/20 mils
Spacing (Strobe-to-Strobe/Data)	5/20 mils	5/25 mils
Total Etch Length for motherboard	7.25 Inches	6.00 inches
Total Etch Length for Daughter Card	1.5 Inches	1.5 Inches

Table 8: Four-Layer Motherboard Stackup Information

	AGP 4X Mode	AGP 8X Mode
Material	FR-4	FR-4
Impedance	60 Ω \pm 15%	60 Ω (microstrip), 56 Ω (stripline) \pm 10%
Board Thickness	62 mils \pm 10%	62 mils \pm 10%
Component Layer Trace Width	5 mils	5 mils
Breakout	5/5 spacing for BGA breakout and for routing in and out of connector	5/5 spacing from BGA breakout (500 mils) and for routing in and out of connector (200 mils).
Spacing (Data-to-Data)	5/15 mils	5/20 mils
Spacing (Strobe-to-Strobe/Data)	5/20 mils	5/25 mils
Total Etch Length for motherboard	7.25 Inches	6.00 inches
Total Etch Length for Daughter Card	1.5 Inches	1.5 Inches

1.4 Receiver

At the receiver, there are three parameters that affect setup and hold timing requirements.

- Intrinsic setup and hold times required by latches at the pad ring: When designing the data receiver latches (flip flops), attention must be given to trading off a faster setup and hold time vs. the clock to output time of the latch.
- Effect of common mode noise between strobe and data, if the strobes are differentially received, as well as noise differences between the transmitter, receiver and the voltage reference (Vref). Refer to Section 1.4.1 and 1.4.1a for more details.
- Spatial skews between the data path (output of the input receiver to input latch) and the strobe path (output of the input receiver with strobe distribution to input latch) as shown in Figure 16 between path A and B respectively. Since Strobe and Strobe# need to be distributed to 16 bits of data, this skew is primarily due to the extra loading on the strobes and the interconnect (RC) mismatch between the data and strobe lines. Careful consideration needs to be taken in the pad ring floorplan

to route the strobe and strobe# to the latches. Minimize by design the skew between the data path and the strobe path with either active or passive components on the data path to account for strobe distribution. An example in Figure 16 shows one possible receiver configuration, where strobe signals are received differentially. In non-differential AGP 4X mode, or AGP 8X, the strobes are being compared with the reference voltage.

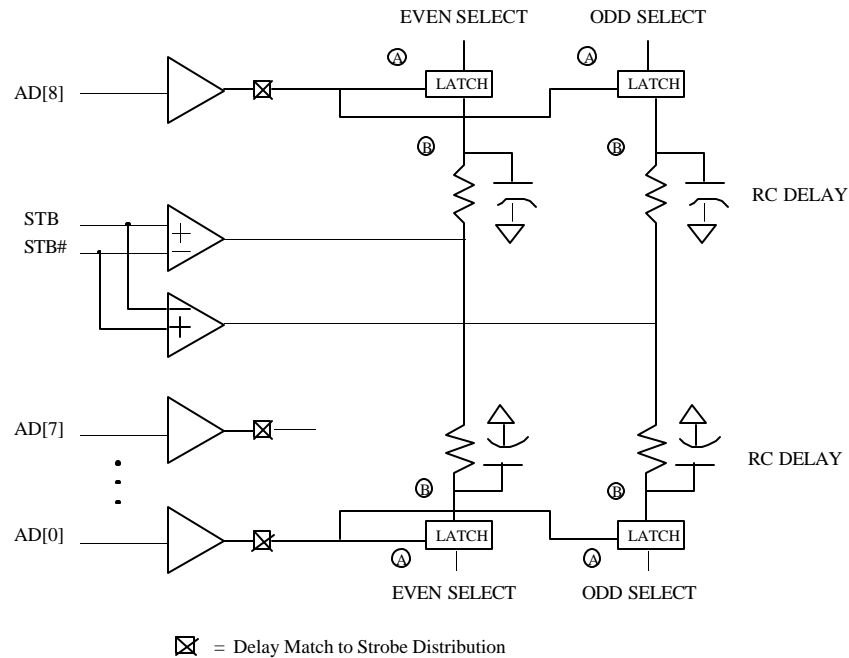


Figure 16: Matching Between Strobe and Data at Receiver for AGP 4X

1.4.1 Input Receiver Considerations – AGP 2X and AGP 4X

Input sense levels can have a significant impact on signal skews. Getting the best timing margins for the AGP 2X 133 MHz and AGP 4X 266 MHz transfer mode requires minimizing signal skews by controlling the input sense levels. It is recommended that designs use differential buffers as input buffers with a reference sense level, Vref, to control skew. Vref has been selected to track the interface voltage. (Vref is 0.4Vddq for 3.3 volt signaling and 0.5Vddq for 1.5 volt signaling.)

Timing uncertainty (skew) is caused at the receiver by the uncertainty of the input buffer sense level over the rise or fall time of the input signal. AGP 2X mode allows the rise/fall rate to be between 1.5 and 4.0 Volts/ns. The 1.2 volt ($V_{IH} - V_{IL}$) TTL input sense range would allow as much as 0.8 ns skew in the sensing of a signal. For AGP 2X mode, the input sense range is $(0.5 \cdot V_{ddq} - 0.3 \cdot V_{ddq})$. At $V_{ddq} = 3.3V$, the uncertainty range is about 0.66V and the maximum skew is approximately 0.44 ns. AGP 4X mode allows the rise/fall rate to be between 1.4 and 3.0 Volts/ns. For AGP 4X mode, the input sense range is $[(V_{ref} + 200 \text{ mV}) - (V_{ref} - 200 \text{ mV})]$. At $V_{ddq} = 1.5V$, the uncertainty range is about 0.40 V, and the maximum skew is approximately 150 ps. A reduced input sense level uncertainty requires a tighter sensing scheme with a sense level common between AGP interfaces. A V_{ref} signal has been defined to provide this input sense level reference. It has been defined to be V_{ddq} dependent so it can be provided from a simple resistor divider. See Section 1.4.4 for more details on V_{ref} characteristics. A differential input buffer can be used to provide accurate switching at the V_{ref} level. The buffer should have a combined input sensitivity and input offset voltage range of less than $\pm 100 \text{ mV}$ for AGP 2X and AGP 4X. The smaller this range can be made, the smaller the timing skew will be between inputs.




Figure 17: Differential Input Buffer and Required Transfer Characteristic

Although typical CMOS input buffers can be used, it is more difficult to make their input sense levels and sense range as stable and narrow as it can be with a differential buffer. Also, the input sense levels of typical input buffers are more susceptible to ground and power noise. A differential input buffer with an external voltage reference is less affected by on-chip noise. The buffer requires high power supply rejection ratio (PSRR) and common mode rejection ratio (CMRR). This can be achieved by using a multistage input or folded cascode receiver.




Figure 18: MultiStage Input Receiver

This input buffer scheme has been adopted to minimize skew. The differential input buffer should be designed to have the same propagation delay for rising and falling transitions so that it is not a significant source of signal skew itself.

The on-chip Vref signal is susceptible to noise from sources such as adjacent data, clock, power, and ground signals. Care should be taken to shield Vref with Vss. See Figure 19 for an example of a three metal layer technology implementation.




Figure 19: On-Chip Signal Metal Layer Cross Section

1.4.1.1 Input Receiver Considerations – AGP 8X

This section describes the impact of noise and Vref tolerances on receiver timings, as they pertain to the AGP 8X mode of operation. The AGP3.0 Specification, Revision 1.0 states that a Vref voltage may be generated by a voltage divider, consisting of $\pm 1\%$ tolerance resistors. When supply variation of $\pm 5\%$ is factored into the equation, the voltage reference can vary from 330 mV to 374 mV. However, it's important to simulate the receiver with reference voltages above and below these numbers to ensure the circuit's functionality. Be sure to verify the noise margin.

Figure 20 shows an example of a waveform, which exhibits a ledge due to imperfect termination.





Figure 20: Ledge Caused by Imperfect Termination

In order for the receiver to trigger correctly, the ledge must not cross the reference. However, even if the ledge doesn't cross the reference, the receiver timing may be severely affected by the limited amount of overdrive level (the gap between the reference and the input signal). If the gap is small, the receiver propagation delay may increase significantly, depending on the design, and impact the setup/hold margins. For example, if the height of the ledge is 170mv, and the Vref is at 330mv, then the receiver has 160mv of overdrive. Even if system simulation shows that there is plenty of headroom to the Vref, the actual headroom may be reduced by the noise difference between the driving chip, the receiving chip, and the receiving reference. The following picture illustrates such a scenario:



$$V_{in}=(V_{txn}+V_{rxn})/2$$

Figure 21: System Simulation

Voltage sources V_{txn} and V_{rxn} represent the transmitter and receiver power/ground delivery network noise (ground bounce). Both noise sources are referenced to an ideal ground. While the motherboard ground is not ideal, it may be an acceptable representation of the ideal ground, since the voltage reference, the transmitting chip, and the receiving chip are eventually referenced to the motherboard. The receiving chip or the transmitting chip (but not both) may be located on an AGP card, which will have additional noise relative to the motherboard. All of this noise should be added into V_{rxn} or V_{txn} . The total noise appearing at the receiver is $V_{in} = (V_{txn} + V_{rxn})/2$. If the value of V_{in} exceeds $V_{ref} - V_{ledge}$, then the system is broken, since the voltage noise will cause the signal to exceed the reference. Even if the signal doesn't exceed the reference, the system may still be impacted by slowing down the response of the strobe or data receiver to the falling edge of the pulse with a ledge. In the case where lowest V_{ref} is 330mv and the V_{ledge} is 160mv, the V_{in} must be significantly (at least 50mv) less than 170mv.

The designer must make sure that the power delivery noise of the system (either the transmitter on the card or the receiver on the card), relative to the motherboard, doesn't exceed 100mv. In this case, the receiver still has $330\text{mv}-160\text{mv}-(100\text{mv}+100\text{mv})/2=70\text{mv}$ of overdrive margin.

Section 2.3.2.10, Signal Quality gives input buffer simulation methods and waveforms to check if the buffer meets the signal integrity requirements. This section also gives a recommended input buffer characterization method.

1.4.2 Termination on Unused Signals (Pull-ups / Pull-downs): AGP 2X/AGP 4X

Unused signals need to be tied to their inactive or no operation state to reduce component noise, power consumption, or possible system malfunction. Signals that may be affected are **SB_STB/SB_STB#** and **AD_STB[0::1]/AD_STB[0::1]#**. Other signals may be included depending on the particular AGP implementation. These signals can be left floating only if their receivers can be disabled, and the component will not be damaged if these pins are floating. For example, the **SBA[0::7]** signals may not be used in a particular design implementation. If these inputs cannot be reliably disabled in silicon, these signals may be pulled up to Vddq to prevent the inputs from floating. If these signals are pulled up, it is

recommended that the stub length to the pull-up be as close to zero length as possible to preserve signal quality and minimize signal skew.

- Signals should be disabled with a pull-up resistor to Vddq (not Vcc). The value should be $8.2\text{ k}\Omega \pm 30\%$.
- Pull-ups must be placed on the motherboard (or internal to the AGP target).
- Internal pull-ups may be disabled under system control to save power in a configuration where they are not needed.
- No inputs to an AGP compliant device should be tied to a voltage greater than the Vddq supply voltage.

For AGP 2X/4X Motherboards, Interconnect stubs for SB_STB/SB-STB#, AD_STB[0::1]/AD_STB[0::1]# should be minimized and kept between 0.0 and 0.250 inches long from the trace to the resistor. The stub should be in-line. For an example, refer to Figure 22. This Figure shows how the stubs can be minimized when the strobcs are routed on the bottom layer while the SMT components are on the top layer.

The stub consideration should be applied to any source synchronous signal that needs a pull-up or pull-down.




Figure 22: Strobe Routing to Termination Resistors

1.4.2.1 Termination of Unused Signals (Pull-ups / Pull-downs) – AGP8X

Designs capable of supporting AGP 8X signaling and AGP 2X or 4X signaling need to integrate the pull-ups and pull-downs needed in 2X and 4X modes on die to avoid the signal integrity issues with stubs and extra board vias caused by external termination resistors. Unused signals should be pulled to their inactive states.

1.4.3 Signal Overshoots and Undershoots – AGP2X/AGP4X

Output buffers are specified relative to the PCB impedance range of 50-80 Ω for AGP 2X and 54-66 Ω for AGP 4X. Designers should consider using compensated buffers to control slew rate, reduce SSO noise, and improve ISI behavior.

The primary cause of signal overshoot and undershoot is the output driver strength. The output driver is driving the line with a lower impedance than the characteristic impedance of the board trace. In a point-to-point topology like AGP, this causes an initial voltage ledge at the driver greater than half swing and a full reflection at the receiver that is beyond the rail. The stronger the driver and the greater the mismatch between the driver and the line, the greater the magnitude of the overshoot. Due to the separate pull-up and pull-down devices in the driver, the magnitudes of the overshoot and undershoot can be independently affected by the driver strength.

1.4.3.1. Signal Overshoots and Undershoots – AGP8X

The PCB target impedance range for AGP8x is 56ohm. A $\pm 10\%$ variation due to PCB making process is allowed. The termination resistor is valued at 50 ohm. This termination provides a significant reduction in overshoot and undershoot. However, impedance differences between motherboard and termination, as well as the AGP connector, will cause reflections. In AGP 8X mode no voltage spikes of >350 mV above Voh max level, or >300 mV below Vss are expected. The highest ledge observed in some system simulations was around 170mv from ground and about that much from the Vswing level

1.4.4 Voltage / Loading Characteristics of Vref

1.4.4.1 Vref Characteristics for AGP 2X Mode

Vref is a DC voltage reference signal used to set the input sense level on the AGP bus. It is set between $0.39 \times V_{ddq}$ and $0.41 \times V_{ddq}$ Volts. It can be generated on the AGP compliant device or provided by the system using a simple resistor divider from Vddq. Vref can be generated at each device or one source can be used by both devices on the bus if they are both down on the motherboard. If it is supplied to a component externally, it must stay within this range with a $\pm 10 \mu A$ DC current load from every device using Vref. Since noise may be generated from other signals coupling to Vref, proper bypassing must be provided.

The system can provide Vref from a simple resistor divider (Figure 23). The resistor ratio R1/R2 must be exactly 2/3. Table 5 shows some sample values of R1 and R2 for resistor tolerances of 1% and 2% with one device load current of $\pm 10 \mu A$. The tighter tolerance resistor has more margin to the Vref specification for IR drop due to the load current and, therefore, can have larger resistance values. If the Vref circuit is used to drive two AGP compliant devices, resistor values no more than half the maximum allowed value must be used due to the doubled load current.




Figure 23: Vref Resistor Divider for AGP 2X Mode

Table 9: AGP 2X Mode Resistance Values

R1 (K Ohm)	R2 (K Ohm)	Tolerance
1.00	1.50	1%
1.10	1.65	1%
1.40	2.10	1%
1.58	2.37	1%
1.62	2.43	1%
1.74	2.61	1%
1.78	2.67	1%
1.96	2.94	1%
2.32 ¹	3.48 ¹	1%
0.10	0.15	2%
0.12	0.18	2%
0.16	0.24	2%
0.18 ¹	0.27 ¹	2%

Note 1: Maximum Standard Resistance Value

The resistor divider should be bypassed to ground near each Vref pin to reduce noise. A standard 0.01 μ F ceramic bypass capacitor for C1 is recommended. Care should be taken in board layout to avoid signal crosstalk from other signal lines. Also, the Vref input pin should be kept away from other switching signal pins to avoid crosstalk in the package. Similar care should also be taken in routing the Vref signal inside the AGP interface circuits to avoid noise pick-up.

Vref can also be generated internal to the AGP interface, so the actual value of Vref may not be visible. In this case, VIL and VIH levels are calculated from and tested relative to Vddq. Similar care has to be taken in this type of design to avoid signals coupling to the internal Vref, and local bypassing may be necessary.

1.4.4.2 Vref Characteristics for AGP 4X Mode

Vref is set between $0.48 \times V_{ddq}$ and $0.521 \times V_{ddq}$ Volts for AGP 4X mode. The common mode relationship between data and Vref at the driving component can be communicated to the receiver by having the Vref generated at the driver (source generated reference). Vref is generated at the driving component using a simple resistor divider network from its Vddq and Vss. Two unidirectional Vref pins are provided in the connector for delivering Vref between the add-in card component and the motherboard component. As an example for outbound writes from the graphics master to the target, the Vref will be delivered using the unidirectional pin on the connector to the target. The Vref is generated using the master component's power supply Vddq and Vss using a simple resistor divider network. Since noise may be generated from other signals coupling to Vref, proper signal spacing between Vref and other signals must be provided. Also, the Vref lines must be matched in length to the data lines in order that the common mode component of the Vref arrives at the receiver at the same time as the data. A length mismatch will phase

shift this common mode signal with respect to the data. A misalignment of the signal can actually reduce the receiver signal margin, so care must be taken when supplying the source generated Vref.

The voltage divider network, shown in Figure 24, consists of an AC and DC element. The DC element is usually made of larger resistors to minimize DC power. The maximum value for these resistors is limited by the leakage current at the Vref input of the device ($\pm 10 \mu\text{A}$). The resistor values of R1 and R2 are equal and must be selected in order to meet this input current load specification. If 1% tolerance resistors are used, any standard value between 80Ω and $2.1 \text{ k}\Omega$ will work. For 2% tolerance resistors, use any standard value between 80Ω and $1.4 \text{ k}\Omega$. The actual value should be selected consistent with the DC current allowances and values common to the kit of materials used on the board.

The AC element is tuned such that the equivalent impedance at Vref approximates the buffer output impedance. This makes the noise bandwidth on Vref the same as on the buffer. For the AC elements, resistors R3 and R4 must be equal in value and 2% or better tolerance. Their value should be selected such that the parallel resistance of R1, R2, R3, and R4 is about 40Ω . Note that if R1 and R2 are selected to be about 80Ω , the AC element of the network is not even needed. For C1 and C2, a low ESR/ESL NPO or X7R type capacitor with a value of 0.5 nF at 4 MHz can be used for the Vref circuit. The Vref resistor divider network must be placed away from critical and noisy signals. Care should be taken in board layout to avoid signal crosstalk from other signal lines.

Also, the Vref input pin should be kept away from other switching signal pins to avoid crosstalk in the package. Similar care should also be taken in routing the Vref signal inside the AGP interface circuits to avoid noise pick-up.

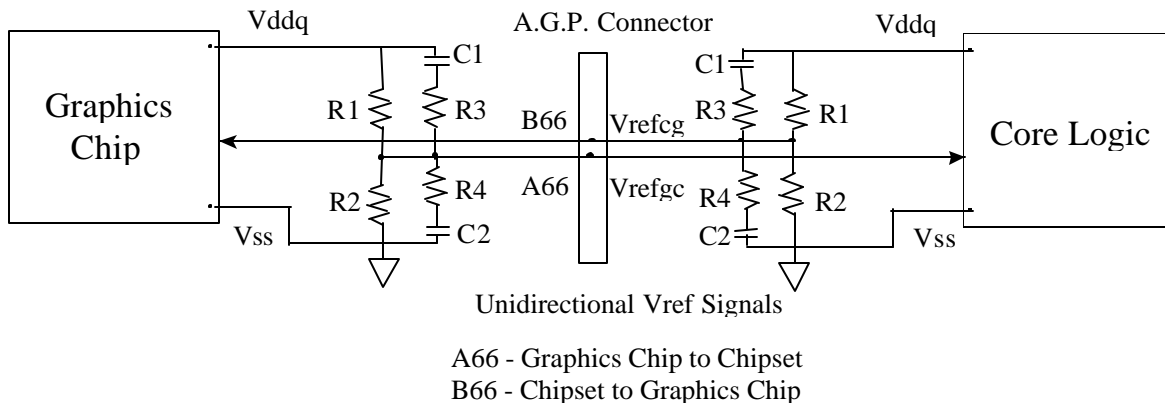


Figure 24: Vref Resistor Divider for AGP 4X Mode

1.4.4.3 Vref Characteristics for AGP 3.0 Mode

As opposed to AGP2.0, Vref for AGP3.0 is highly recommended to be locally generated – for the chipset on the motherboard and for the graphics controller on the add-in card.

The method of generating VREF for and AGP3.0 mode is similar to the AGP2.0; it's a simple voltage divider with a capacitor referencing this voltage to ground (see Figure 26). The value of resistors is arbitrary, as long as the reference is set to 0.35 volts for Vddq at 1.5 volts, as close as standard 1% tolerance resistors allow. The smallest value of the resistors depends on the acceptable current consumption by the divider, as well as the filtering effect of the capacitor that provides the ground reference. The largest value is set by the error created by input leakage currents of the receiver.

The amount of decoupling capacitance used for referring the Vref node to ground depends on the amount of coupling desired at specific frequencies. If motherboard ground is viewed as a source of noise, the designer can calculate the amount of that noise coupling into the Vref node by using the circuit below.




Figure 25: Noise Coupling Circuit

$(s+1/R1 \times R2 \times C)/(s+(R1+R2)/C \times R1 \times R2)$ is the transfer function for the above circuit, which is a high pass function with DC gain at $R1/(R1+R2)$ and an AC gain at infinity equal to one. The $R1$, $R2$ & C values can be chosen based on the desired amount of noise coupling. A small capacitor value will cause poor coupling of V_{ref} to ground, while a high value will produce high coupling. If we chose $R2$ to be 1K, and C to be 0.01uF, then >98% of ground noise of >1MHz will be coupled to V_{ref} . $C=0.01\mu F$ is, therefore, more than sufficient for this purpose.

Since, the voltage references for AGP 2.0 and AGP3.0 are different, it's also necessary to implement a mechanism for switching from one to the other. This is discussed in more detail in the next section.

1.4.4.4 AGP3.0 VREF Generation Using Detect Signal Circuits

Figure 26 shows an example of a voltage divider that can be used to generate V_{REF} from V_{ddq} for an AGP3.0 only system.




Figure 26: Vref Resistor Divider for AGP3.0 Mode

To generate the proper V_{ref} for a universal AGP2.0/3.0 board or add-in card, either of the following

circuits shown in Figure 27 can be used. **MB_DET#** is generated from the motherboard and **GC_DET#** is generated from the add-in card.

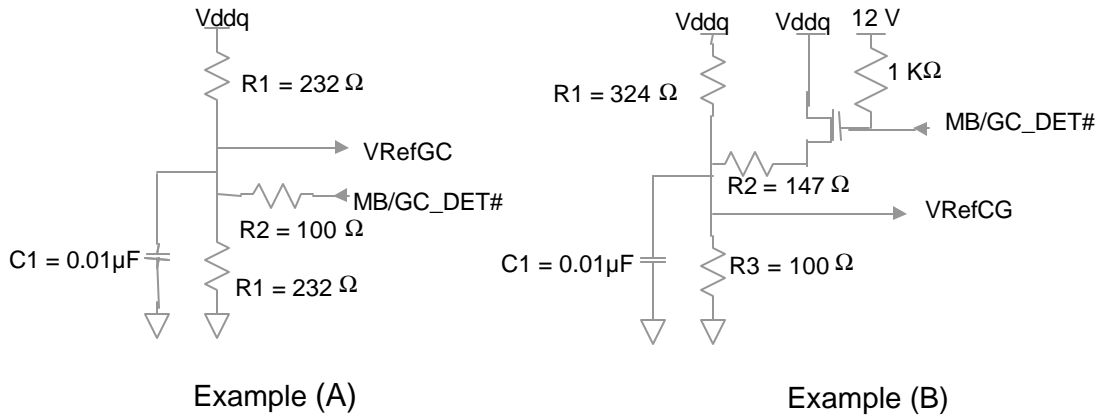


Figure 27: Vref Resistor Divider for All 1.5V Modes

When example A or B is used on the motherboard, then the correct signal to connect is **GC_DET#**. When the circuit (A or B) is used in the add-in card, the correct signal to connect is **MB_DET#**.

For either of the circuits above, the VREF generated by the circuit placed on the add-in card is supplied to pin A66 on the AGP connector (**VREFGC**). The VREF generated at the motherboard level is connected to pin B66 on the AGP connector (**VREFCG**).

Supplying the AGP3.0 VREF to the AGP connector is required. Use of the VREF supplied through the connector is optional and implementation dependent. If the connector-supplied VREF is used, a local bypass capacitor (0.01 μF) to ground should be placed close to the receiving device's VREF pin to reduce noise.

1.4.4.5 Vref Margining

One method of determining the robustness of a system design is to use Vref margining. The resistor divider source of the Vref voltage is replaced with a variable resistor of the same total value. With the system running an AGP interface data intensive application, Vref is varied around its nominal value. Vref should be able to be varied by several tens of millivolts in both the positive and negative direction without creating data corruption or system hangs.

Lack of margin may be due to noise coupling on to the Vref line, or due to signal integrity problems in the data or strobe signals, which can be verified with an oscilloscope. Vref noise may require that the signal be rerouted or that the source be cleaned up. If the source generated Vref is used, the Vref noise may not be correlated to the data noise. The Vref may have to be locally bypassed or the design may have to be switched to a locally generated Vref.

Data and strobe signal integrity issues are usually caused by impedance discontinuities in the signal path or power delivery problems to the drivers. Impedance discontinuities occur when the signal changes transmission media (package to board to connector, etc.) or when the signal changes its return reference plane (changing routing layers). These discontinuities need to be minimized. Signal return layer should be left the same for the whole trace, and traces must not cross power plane splits. Where discontinuities do occur, they need to be bypassed to allow return currents a low impedance path. Power delivery issues can be improved by proper bypassing and power plane layout (see Section 1.5.2).

1.5 AGP Vddq Power Delivery

Power delivery to AGP components, specifically for the Vddq for the graphics controller, entails the following areas of focus on both the motherboard and add-in card: The voltage regulator selection and design, motherboard parasitics, connector parasitics, power plane decoupling, card parasitics, and the transient current load waveform.

The proper understanding of these parameters is key for meeting the voltage target for Vddq which is to be applied at the card side of the connector. The voltage target is 1.5V +/- 5%.

Below, in Figure 28, is a circuit model that serves as a guide or example in showing values for all the above-mentioned focus areas, and helps add perspective towards modeling and simulating the power delivery from a system viewpoint. The various elements or parasitics in the model below will be referred to in following sections and will represent the write case when the graphics controller is driving which, for the purposes of this section, means that the graphics controller is presenting a current load.

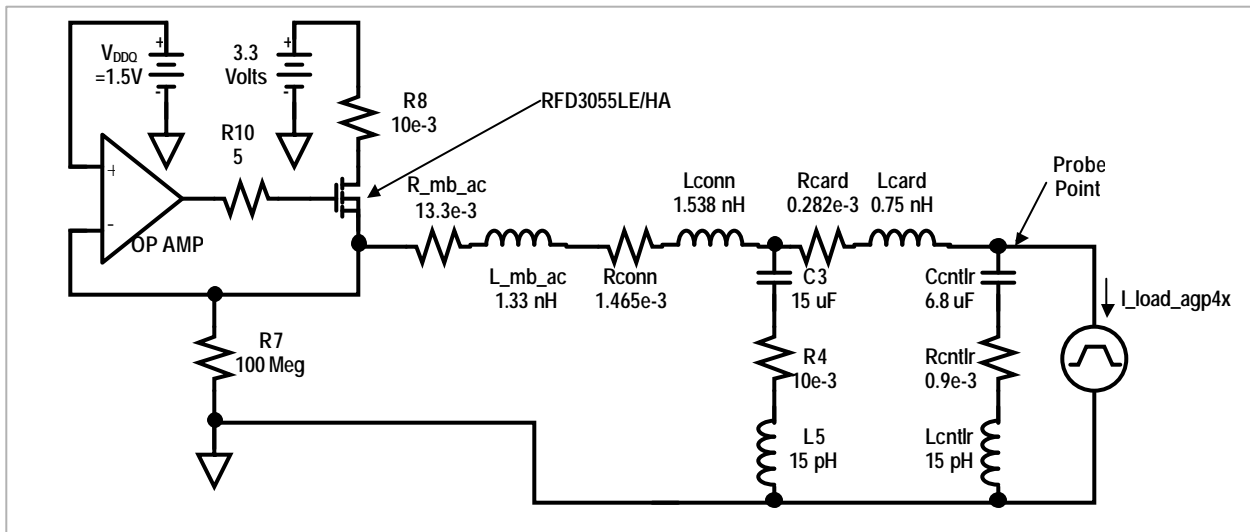


Figure 28: Vddq System Power Delivery Model for Graphics Controller Driving

1.5.1 Vddq Voltage Sources

Power to the Vddq rail may be supplied by a voltage regulator or the system power supply depending on system configuration. In the case of a 3.3V only signaling environment, the system power supply will most likely be used. In the case of a 1.5V (which is only supported in AGP 8x mode), or universal 3.3V/1.5V signaling system, a voltage regulator will most likely be used. In the universal case, the voltage level of the regulator is controlled by the **TYPEDET#** pin which is controlled by the add-in card.

The two main types of regulators that can be used for the Vddq rail are linear and switching. A linear regulator is defined as a voltage control circuit with a pass transistor in the current path. A switching regulator is defined as a voltage control circuit with a switched inductive tank circuit used in the current path. Each type of regulator has associated advantages and disadvantages including: cost, part count, board area, response time, efficiency, and thermal dissipation. Examples of a linear and a switching regulator are shown in Figure 29 and Figure 30, respectively:

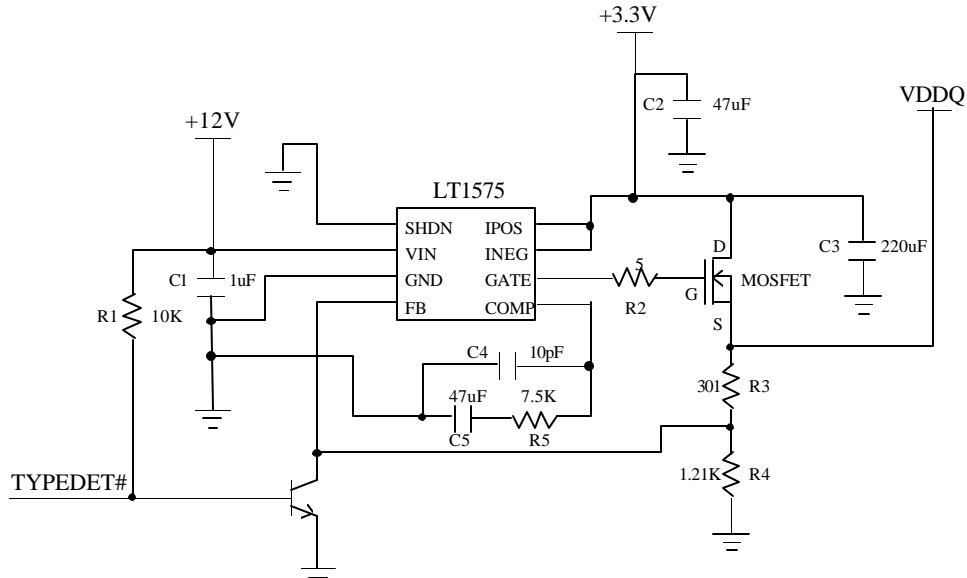


Figure 29: Example of a Linear Regulator Circuit

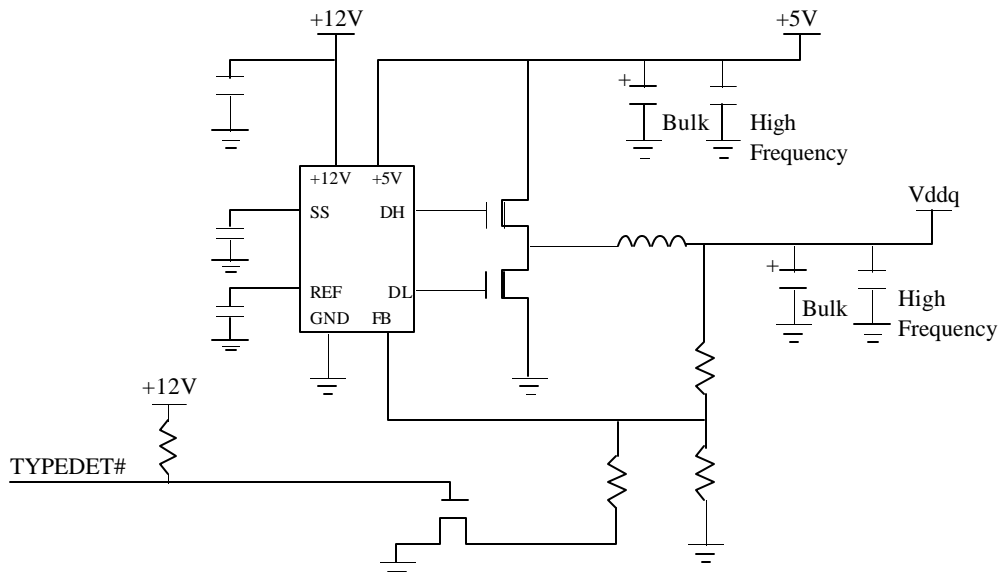


Figure 30: Example of a Switching Regulator Circuit

1.5.2 Motherboard Parasitics

It is important to include the effects of the plane parasitics on the motherboard, since they can contribute to a relatively large inductance and resistance. The potential for significantly high values of these parasitics stems from the Vddq/Vss routing, the number of via holes, and the dimensions of the planes. To capture the effects of holes, irregular geometry and the concerns of current crowding, the use of a 3D field solver is recommended.

In the model shown in Figure 28, the values for “R_mb_ac” and “L_mb_ac” are the AC resistance and inductance of the Vddq and Vss planes under the condition of the graphics controller driving. Another set of parameters can be obtained, by following the same methodology, for the read case when the

MCH/chipset is driving and the graphics controller is receiving.

These values are obtained by use of the Ansoft Maxwell Q3D\$ Extractor and are particular to layout of the Vddq/Vss planes for an implementation. In order to obtain the values used in the model, it is important that the geometry of the Vddq and Vss planes computed in the 3D solver are closely representative of the actual board implementation of the planes. It is also very important to establish the nodes of the current path by specifying source and sink (load) points for each conductor. Frequency and material properties are equally important for including the effects of skin effect and material resistivity which impact the AC resistance.

1.5.3 Connector Parasitics

Besides the effect of the system/motherboard parasitics, the connector also needs to be modeled since the inductance of the connector is typically one of the largest in the entire model. It is important that the inductance effect is captured since that impacts how effective the bulk decoupling capacitors on the motherboard (between the voltage source/regulator and the connector) will be in terms of delivering enough charge / current fast enough to respond to the load demands of the graphics controller. If the connector inductance is large, the motherboard bulk decoupling could be rendered fairly ineffective in terms of responding quickly enough to di/dt slew rate demands of the graphics controller. Therefore, card caps may be the best way to go. This can only be found out by including the parasitic effects of the connector into the model. For the model shown in Figure 28, the values for the connector contact resistance and effective inductance are found in the *AGP Interface Specification, Revision 2.0*

- Initial contact resistance (section 5.6.1.2) = 10mOhm max for any pwr or ground contact before testing
- Effective inductance (section 5.6.3) = 10.5 nH max

Besides knowing the parasitics per pin, it is important to also know the total number of Vddq and Vss pins, therefore the *AGP3.0 Interface Specification, Revision 1.0* (Section 2.8) was referenced to get the information below:

- # of current carrying Vddq pins = 11
- # of current carrying Gnd pins = 18
- # of current carrying Vcc3.3 pins = 8

Impedance for AGP 8X or 1.5 V (covers 4X):

$L_{\text{connector_total_Vddq}} = 10.5 \text{ nH} / 11 = 0.95455 \text{ nH}$

$L_{\text{connector_total_Gnd}} = 10.5 \text{ nH} / 18 = 0.5833 \text{ nH}$

$R_{\text{connector_total_Vddq}} = 10 \text{ mOhm} / 11 = 0.91 \text{ mOhm}$

$R_{\text{connector_total_Gnd}} = 10 \text{ mOhm} / 18 = 0.556 \text{ mOhm}$

$L_{\text{connector_lumped_total}} = 0.95455 \text{ nH} + 0.5833 \text{ nH} = \mathbf{1.538 \text{ nH}}$

$R_{\text{connector_lumped_total}} = 0.91 \text{ mOhm} + 0.556 \text{ mOhm} = \mathbf{1.465 \text{ mOhm}}$

1.5.4 Power Plane Decoupling

Since any regulator or power supply response time is much too slow for the transients present in the AGP signaling environment, capacitive decoupling of the power planes is essential. These transients can cause voltage drop on the power plane until some mechanism for power delivery can supply the needed current. Capacitors are used to supply power during transients until the voltage regulator can respond. Quite often in a system, two types of capacitors are needed to perform adequate decoupling: high frequency and bulk.

High frequency capacitors are small in value (0.001 μF – 1 μF) and have very low equivalent series resistance (ESR) and equivalent series inductance (ESL). Although the capacitance of a high frequency

capacitor is not large, it can deliver energy quickly to a power transient due to its low ESR and ESL. However, since the capacitance is low, it cannot sustain the power for large periods of time.

Bulk capacitors are generally larger (10 μF – 10,000 μF) and have higher ESR and ESL than high frequency capacitors. The higher ESR and ESL make their response time to power transients longer and thus they are less effective at preventing voltage drops during these transient conditions. The larger capacitance of bulk capacitors does, however, allow them to deliver power to the system for longer periods of time than high frequency capacitors. The bulk capacitor can be thought of as a temporary voltage regulator that maintains the rail power and the high frequency capacitor power until the voltage regulator can respond to the power transients.

Typically, in previous designs, there have been high frequency caps very close to the graphics controller and the connector on the add-in card. If the board technology and costs allow, capacitors should be mounted on the backside of the card under the controller package with short direct connections between the power and ground vias supplying the chip. Bulk decoupling has typically been mounted near the chipset interface and near the connector for not only power delivery but also for signaling reasons.

The effectiveness of any bypass capacitance is limited by the inductance and resistance in the bypass capacitor as well as the inductance that is between the capacitor and buffers on the die. This includes board vias, traces, planes and package inductance. Anything that can be done to reduce this inductance will improve the power delivery and associated signaling quality. Using more capacitors reduces inductance by paralleling the paths. The capacitors should be as close to the power pins (package or connector) as manufacturing rules allow. Watch out for narrowing of power delivery paths due to anti-pads in the power plane. Power delivery inductance can be reduced by flooding or fingering V_{cc}/V_{ddq} on a signal plane over the ground plane (or visa versa) to create a closely coupled path where the mutual inductance between the planes reduces the self inductance.

If the decoupling is not sufficient in value or is too far away (too much series inductance), V_{ddq} will droop below the specified minimum value. The driver characteristics are generally not specified in this range and will not deliver good signal quality. The more the V_{ddq} droops (or ground supply rises), the more the timing of the buffer will be pushed out. The output slew rate will also be less. Slow slew rates cause greater timing skew in the interconnect, further degrading system timing. V_{ddq} noise can also be coupled on to the V_{ref} signal causing reduced switching margins and timing skews. Good power delivery is key.

In Figure 28, C3, R4, and L5 can be located. C3, R4, and L5 represent the total effective capacitance and associated parasitics for the high frequency decoupling to be placed near the connector on the card. From our analysis, HF decoupling designs must meet the requirements of C3, R4, and L5 in order to meet the voltage targets. C3 represents the minimum capacitance required (15 μF), R4 represents the maximum allowable ESR (10 $\text{m}\Omega$), and L5 represents the maximum allowable ESL (15 pH). Maximum and minimum totals are shown on purpose to give those following this guide an opportunity to optimally pick decoupling capacitors that meet their needs depending on the constraints of their design environment (such as cost, size, real-estate). No matter what capacitor is selected, the effective high frequency capacitance to be placed on the card near the connector must be $\geq 15 \mu\text{F}$, the associated total ESR must be $\leq 10 \text{ m}\Omega$, and the total ESL must be $\leq 15 \text{ pH}$ in order to meet voltage targets.

As an example of a type of capacitor to select that will meet the maximum and minimum capacitance and associated parasitic requirements stated above, while still meeting the voltage spec, see Table 10 below. The capacitance information in Table 9 is representative of the write condition which is when the graphics controller is driving.

Table 11 shows additional high frequency capacitors that are placed near the connector, but this time, on the motherboard side. These caps are useful during the write condition which is when the graphics controller is driving and primarily function to assist in the signal integrity areas of return path and via stitching. Low ESL and ESR capacitors are also needed here.

Table 12 shows that two additional capacitor stuffing spots for bulk decoupling are planned for on the motherboard. The analysis in Figure 28 does not require bulk decoupling for the write condition in which the graphics controller is driving. However, since this decoupling has been used in previous designs, it is recommended to allow stuffing spaces, as stated in Table 11, in case the bulk decoupling is needed for power delivery.

Table 13 shows high frequency decoupling caps placed near the MCH/Chipset to assist in the signal integrity areas of via stitching and return paths for the read condition (MCH/chipset is driving). These

caps will assist in power delivery. Again, low ESL and ESR capacitors are needed here.

Table 10: Example HF Capacitors Placed Near the Connector on the AGP Card

Vendor	Pkg	Qty	C/ea.	ESR/ea.	ESL/ea.
AVX	Pkg. 1206, 8 terminal	8	2.2 uF	3.88 mΩ	110 pH

Table 11: Example HF Capacitors Near the Connector on the Motherboard

Vendor	Pkg	Qty	C/ea.
AVX, MuRata, TDK, etc.	0603	9	0.1 uF

Table 12: Example Bulk Capacitor Stuffing Locations Near the Connector on the Motherboard

Vendor	Pkg	Qty	C/ea.
Sanyo, Nichicon, Panasonic, etc.	Electrolytic (cylinder shaped)	2	100 uF

Table 13: Example HF Capacitors Near the MCH on the Motherboard

Vendor	Pkg	Qty	C/ea.
AVX, MuRata, TDK, etc.	0603	10	0.1 uF

1.5.5 Card Parasitics

Just as the motherboard plane parasitics need to be considered, so do the card parasitics. These card parasitics lie between the connector card edge and the graphics controller interface, and can also be substantial. Once again, depending on the geometry of the planes and complexity of the plane shapes, the card parasitics may or may not need to be modeled. If they are not modeled, then a hand calculation should be used. If the add-in card parasitics are negligible, then most likely the card capacitors placed close to the card connector edge would suffice for meeting voltage targets. However, if the parasitics are large (on par with or greater than the total ESL and total ESR of the card caps close to the connector), then additional card caps will likely be needed as close as possible to the graphics controller interface.

Below, in Table 14, is an example of the type and quantity of card caps placed near the graphics core on the AGP card. The equivalent capacitance and associated total parasitics can be seen in Figure 1 and have the reference names of “Ccntl,” “Lcntl,” and “Rcntl.” These caps were added so as to meet the same voltage targets that are supposed to be applied at the connector on the card side. Note that these high frequency capacitors shown in Table 13 below are not requirements, but simply serve as an example of what kind of high frequency caps would be needed on the card near the graphics controller in order to meet the same voltage target that is applied to the card side of the connector under a write condition (the graphics controller is driving).

Table 14: Example HF Card Capacitors Placed Near the Graphics Controller

Vendor	Pkg	Qty	C/ea.	ESR/ea.	ESL/ea.
MuRata	0805, 8 terminal	10	0.68 μ F	9 m Ω	150 pH

1.5.6 Transient Current Load Waveform

The transient current load waveform used in the system simulation and shown in Figure 28, “I_load_agp4x,” should be the maximum load. In order to achieve the maximum load from both a static and transient perspective the following parameters should be utilized in the respective tolerance ranges: Lowest Z_o (MB and card transmission line impedance), highest buffer dV/dT , lowest driver impedance, lowest receiver impedance, and the worst case number of lines switching. These parameters yield a worst-case transient current load waveform.

Since all power delivery designs need to be backward compatible (should satisfy AGP 8X and AGP 4X), the worst case load should be identified. It was found that the AGP 4X designs present the worst case transient load. Therefore, Figure 28 shows the load specific to AGP 4X. The AGP 4X current load waveform is being shown below in Figure 31.




Figure 31: Worst Case Current Load Waveform (From AGP 4X) To Be Used In VDDQ Power Delivery Simulations

1.5.7 Connector AC Signal Decoupling Requirements

The decoupling capacitor recommendations for the AGP connector are intended to address A.C. signaling issues and not power delivery issues. The main reason for not addressing power delivery issues with the motherboard connector decoupling is due to the connector inductance and the distance the capacitors are from the graphics device. These two factors negate much of the usefulness of the connector decoupling on the motherboard for power delivery purposes.

The following are recommendations for decoupling at the AGP connector on the motherboard:

- Vcc3.3: Four (three for a 1.5V connector) 0.01 μ F or larger, low ESL capacitors. Each capacitor should be placed as close as possible to a Vcc3.3 pair of pins on the connector.
- Vddq: Six 0.01 μ F or larger, low ESL capacitors. Each capacitor should be placed as close as possible to a Vddq pair of pins on the connector.

- +5 V: One 0.01 μ F or larger, low ESL capacitor placed as close as possible to the +5 V connector pins.
- +12 V: One 0.01 μ F or larger, low ESL capacitor placed as close as possible to the +12 V connector pin.
- 3.3VAUX: Two (One for 1.5V connector) 0.01 μ F or larger, low ESL capacitors placed as close as possible to the 3.3VAUX connector pin(s).

All power and ground pins must be connected on the motherboard to guarantee power delivery and proper AC signal return paths.

All ground pins on the add-in card need to be connected for proper grounding and AC return paths. All Vddq pins need to be connected on the add-in card to the local Vddq power plane. The add-in card does not have to use all of the Vcc3.3 power pins if it does not need the power delivery. However, all unused power pins must be bypassed to ground close to the connector with a 0.01 μ F or larger, low ESR/ESL capacitor to provide good AC coupling to the adjacent signaling pins.

1.6 Motherboard and Expansion Card PCB Layout Considerations

This section describes layout and routing guidelines to ensure a robust AGP interface design. See Chapter 2 for more detail on these topics and specific layout requirements. These guidelines ensure that the AGP specifications can be met and that the motherboard and daughter card will operate together; however, it is not a guarantee that the AGP specifications will be met. The system designer must analyze and simulate the AGP system to ensure that all specifications are met.

1.6.1 Motherboard Layout Recommendations

The following figures (Figure 32, Figure 33, and Figure 34) are examples of AGP layouts for ATX, LPX, and Micro-ATX form factor motherboards. The goal of the placement is to allow the routing of the AGP bus that minimizes trace length, vias, and interference with other signals and buses.

The pinout of the core chipset should allow the AGP bus to flow to the AGP connector or graphics device with a minimum trace length and signal crossing. The AGP strobe signals must be grouped with their associated data signals. It is recommended that the strobe be centered within its group to minimize the signal to strobe skew. The strobe and data groupings for AGP-4X and AGP-8X are shown in Table 15.

Table 15: Strobe and Data Groupings for AGP4X and AGP8X

AGP 2X / 4X		AGP 8X	
<i>Strobes</i>	<i>Data Group Association</i>	<i>Strobes</i>	<i>Data Group Association</i>
AD_STB0* AD_STB0#*	AD[15::0] C/BE[1::0]#	AD_STBF0 AD_STBS0	AD[15::0] C#/BE[1::0]
AD_STB1* AD_STB1#*	AD[31::16] C/BE[3::2]#	AD_STBF1 AD_STBS1	AD[31::16] C#/BE[3::2] DBI_HI, DBI_LO
SB_STB* SB_STB#*	SBA[7::0]	SB_STBF SB_STBS	SBA#[7::0]

* The complement strobe signals are available in the AGP 4X and 8X modes only.

Signal referencing discontinuities are a key concern when routing the motherboard. A discontinuity will cause an impedance change and reflection on the signal than can severely harm signal quality if not

eliminated or controlled. A signal is ground or power referenced when the signal is most closely coupled to that power plane.

The traces should be as short and direct as route length matching rules allow. Avoid changing the power plane reference during routing. Any change in reference plane will need to be bypassed by capacitors placed as close to the vias as possible. Signals may have a power or ground reference inside the package, and it is beneficial to continue that reference outside the package to avoid a signal reference discontinuity. Check with the device vendor for this information. Breakout routing from a BGA package may require taking some signals out on top of the board while sending others through the board to the backside to escape. The proper signal reference should be restored as quickly as possible. The segment that is not properly referenced may need capacitors nearby to allow the signal return currents to switch power planes.




Figure 32: AGP PCB Layout for ATX Form Factor

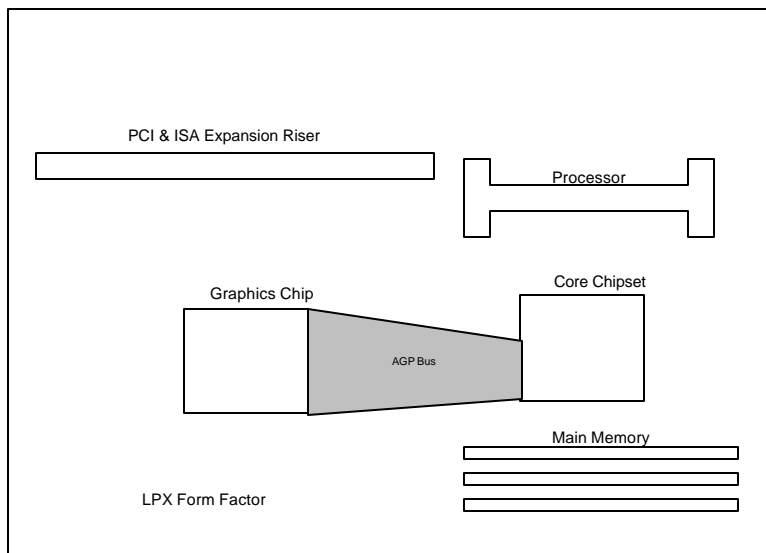


Figure 33: AGP PCB Layout for LPX Form Factor

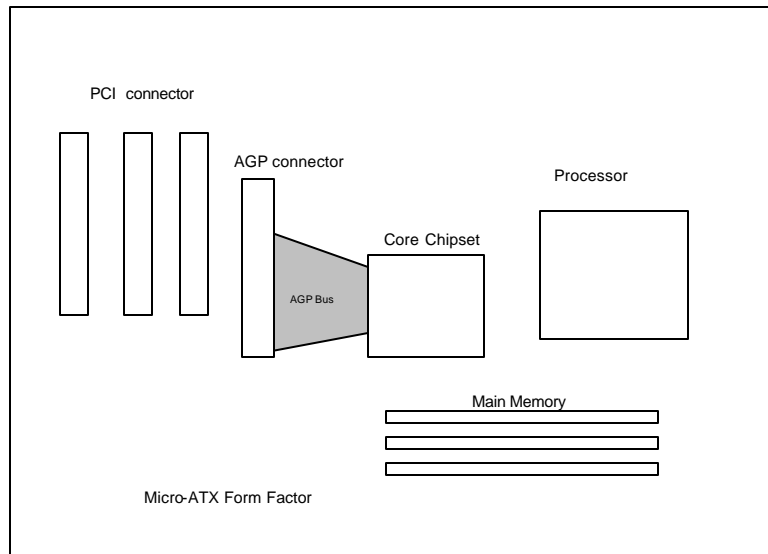


Figure 34: AGP PCB Layout For Micro-ATX Form Factor

The power plane design should be considered during the entire design process and not left to the end. The planes should be cut so that number of “neckdown” (narrowest plane width) points are minimized.

The maximum sustained current through the AGP connector on the 3.3V rail is 6A. Therefore, sustained minimum total width should be no less than 400mils (assuming 1oz copper planes). Brief neckdown points can occur but they will add to voltage drop, and if too narrow, may overheat or burst the plane. At no point should the sum of all 3.3V plane path widths to the AGP connector be less than 150 mils total. Less plane width will not be able to dissipate localized heat generation. Similar power plane width requirements apply to the graphics and chipset components, scaled to their current requirements.

Vddq max sustained current is significantly less at about 1.5A for both the connector AGP components. Vddq minimum total plane width should be 50 mils. At no point should the sum of all Vddq plane path widths be less than 25 mils for the same reasons as above.

Power planes are typically narrowest when breaking out from the power pins for the connector and chip components. The power plane cuts must be designed such that anti-pad layers from surrounding pins do not choke off the plane. For example, Figure 35 shows the power layer of the AGP connector section of a motherboard. Signals that aren't AGP voltages are voided out on the power layer and are denoted as the large black circles (anti-pads). The small circles with crosses are the pins that are tied to the plane. The black line that snakes its way through the voids is the power plane cut. The cut separates the northeast portion of the drawing from the southwest portion. The filled in area between the anti-pads is the power plane. Note how the pin tie on the right has two paths between the anti-pads. Each path is approximately 20 mils wide. This gives approximately 40 mils to this pin. A 40-mil breakout per pin is sufficient to meet the 150 minimum total plane width (6 pins x 40 mils/pin = 240 mils). Individual breakouts less than 25 mils should be scrutinized to try to obtain a wider power plane path.




Figure 35: Power Plane Etch and Voids

General rules for plane cut design are as follows:

- Resistance doubles as length doubles
- Resistance doubles as width halves
- All other variables equal, use of an outer layer doubles effective maximum current, but halves effective temperature rise and cross-sectional area for a given current.
- All other variables equal, temperature rises linearly with current and rises as the square of the change in trace width or thickness.

External power traces for 3.3VAUX and +12V are commonplace in today's designs. External power traces should be kept as short as possible (under 10"). Trace widths minimum of 40 mils for 3.3VAUX and 25 mils for +12V are desirable. These may neck down to 25 and 20 mils for brief lengths to route through tight areas without any temperature rise, but the DC drop will increase slightly.

1.6.2 Expansion Card Layout Recommendations

All AGP signals on the graphics chip should be located so that they are in the same order as the pinout of the AGP connector. This assumes that the component side of the add-in card is as defined in the mechanical section of the *AGP Interface Specification, Revision 2.0*. This alignment will minimize the overall trace lengths and aid in matching the trace lengths within its groups.

The strobe signals must be grouped with their associated data group as shown in Table 15.

The discussion of signal reference discontinuity on the motherboard applies to expansion cards as well. However, the short trace lengths and the requirements to meet the connector pinout make it difficult to maintain one signal reference. In this case, it is best to keep routes short and direct, consistent with the above matching requirements. The number of vias on any signal should be minimized. Proper bypassing of Vddq at the signaling interface and at the connector is critical to providing a low impedance path for the signal return currents that have to change reference planes.

If the pinout is optimized for the expansion card connector card, the AGP interconnect will be reversed with respect to the chipset when the graphics controller is mounted directly on the motherboard. Sufficient space should exist on the motherboard to route the bus and meet the AGP requirements in this case.

1.6.3 Board Impedance

The motherboard and add-in card impedance should be controlled to minimize the impact of any mismatch between the motherboard and the add-in card. The PCB should be fabricated using FR-4 with an overall board thickness of 62 mils $\pm 10\%$. The outer layers should be made using $\frac{1}{2}$ oz Cu, and the inner layers should be made with 1 oz Cu. The AGP 2X board impedance is targeted at $65\ \Omega \pm 15\ \Omega$ for a four layer PCB stackup. An impedance of $60\ \Omega \pm 10\%$ is recommended for AGP 4X and 8X. The outer layer dielectric thickness is thinner for AGP 4X and 8X thus lowering the impedance. Refer to Section 1.3 for AGP 2X, AGP 4X, and AGP 8X board layer stackup.

Lower trace impedance will reduce signal edge rates, decrease over/undershoot, and have less crosstalk than higher trace impedance.

1.6.4 Pull-ups and Termination


All pull-up resistors required by the AGP interface specification must be implemented on the motherboard. The AGP interface was designed to not require any special signal termination. It is up to the designer to simulate the signals to ensure that signal quality requirements are met. Pull-up requirements are dependent on the chipset used. Sometimes A.G.P Core logic vendors integrate some pull-up/down resistors. Check with your AGP Core logic vendor for details. AGP 8X requires integrated termination for designs that also support AGP 2X and 4X.

1.6.5 Board Routing Recommendations

AGP signals must be carefully routed on the motherboard and graphics card to meet the timing and signal quality requirements of the interface specification. The following are some general guidelines that should be followed. Trace lengths included in this section are guidelines only. It is recommended that the board designer simulate the routes to verify that the specification is met.

1.6.5.1 Maximum Trace Length Requirements

This section describes how the maximum trace length is segmented between the motherboard and the expansion card. For the common clock mode, the maximum flight time allowed for the bus is 2.5 ns (2.1 ns for AGP 4X). The flight path timing budget for the motherboard is 1.65 ns and for the expansion card is 0.7 ns. Thus, the maximum trace length on the motherboard is 7.5 inches and on the add-in card is 3 inches. Figure 36 shows how the maximum trace length is segmented between the motherboard and the expansion card for the common clock mode. AGP 2X, AGP 4X and AGP 8X source synchronous signals are limited by flight time skew; therefore, the budgets may not be able to reach those specified for the common clock mode.



A = Expansion Card Timing Budget = 0.70 ns = ~ 3.0 inches

Figure 36: Common Clock Timing Budget Between the Graphics Chip, AGP, and Core Chipset

The maximum trace length is also dependent on the effect of crosstalk on signal skew. When the effect of crosstalk on signal skew is taken into account, the maximum length of the motherboard and add-in card traces is reduced. The amount of crosstalk depends on trace length and trace spacing.

For AGP 2X mode with a 1:2 trace spacing, the distance between traces (air gap) being twice the trace width, the maximum trace length on the motherboard remains at 9.5 inches. For a 1:1 trace spacing, distance between traces (air gap) being equal to the trace width, the maximum trace length on the motherboard reduces to 4.5 inches. Add-in cards should be routed with a 1:2 trace spacing and are restricted to a maximum length of 3.0 inches.

For AGP 4X add-in cards and motherboards, the signals can be routed out of the graphics controller with a 1:1 trace spacing. However, after breaking out from the pins, the source synchronous traces (data and strobe signals) on the motherboard can be six inches long if a spacing of >3 times the maximum dielectric thickness between the layer that the traces are routed on and the nearest voltage plane (V_{ss} , V_{ddq} , etc.) is used. If the space is >3 times the dielectric thickness, then the trace can be up to 7.5 inches. The add-in card is required to use a 3X spacing and is limited to 1.5 inches. In any case, the maximum flight time for the entire trace is 2.1 ns (AGP 4X only). It is also recommended that the spacing between the strobe pairs be between three or four times the maximum dielectric thickness. The strobes should also be spaced to other traces by a minimum of five times the maximum dielectric thickness. The signal line lengths should be matched as close as possible from the breakout area, as shown by the dotted line in Figure 37, to the add-in card edge fingers. The same holds true from the core chipset to the connector.

For AGP 8X add-in cards and motherboards, the breakout for AGP 4X applies. As soon as (tight spacing limited to 500 mil for breakout) a greater spacing is possible, the data signals should have a spacing of 4 times the maximum dielectric thickness between the signal layer and the voltage plane to which it references. For the add-in card, the trace length matching should be done from the dotted line to silicon pad. The same should be done from the AGP connector to the core logic silicon pad. Consult your core logic and graphic controller vendors for the appropriate package lengths.

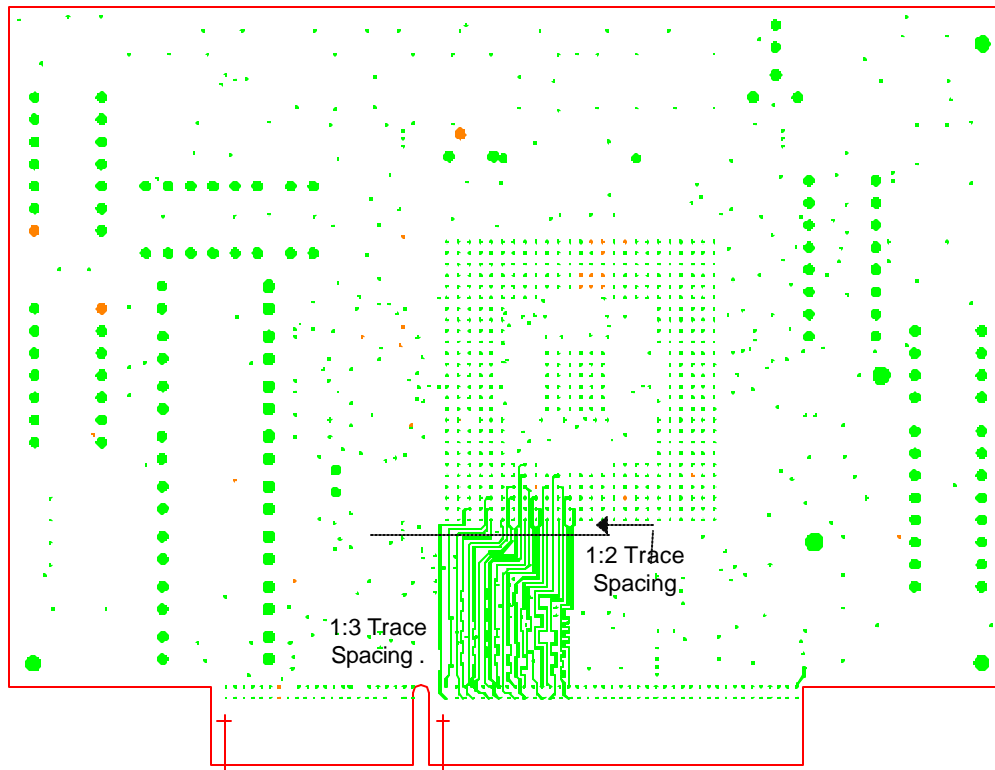


Figure 37: Trace Length Routing Ratio

It is recommended to keep the routes as short as possible to maximize timing margin and signal quality.

1.6.5.2 Trace Length Mismatch Requirements

The trace lengths for signals within a group must be matched to their respective strobe to meet the maximum mismatch requirement given in the AGP Interface Specification. Of the maximum 0.7 ns, the motherboard is allotted 0.5 ns and the add-in card is allotted 0.2 ns. For AGP 4X mode, the motherboard is allotted 25 ps, and the add-in card is allotted 25 ps. For AGP 8X mode, the motherboard is allotted around 10ps, and the add-in card is allotted around 10 ps. A mismatch of 10 ps translates to ± 25 mil for each end of the interconnect and a total of ± 50 mils for the overall matching.

When the effects of varying board impedances and crosstalk are considered, the following trace length constraints are recommended:

- **Add-in Card:** For AGP 2X, the data lines should be kept to within ± 0.5 inches of their respective strobe. For example, if the strobe is at 2.5 inches, the data lines can be 2.0 to 3.0 inches in length. For AGP 4X, the data lines should be kept to within ± 0.25 inches of their respective strobes. For AGP 8X, ± 0.025 inches from their respective strobe is required.
- **Motherboard:** For AGP 2X mode, the data lines should be kept at +0.0, -0.5 inches from their respective strobe. For example, if the strobe is at 9.0 inches, the data line can be from 8.5 to 9.0 inches in length. The data lines should be kept within ± 0.125 inches of their respective strobe for AGP 4X and ± 0.025 inches from their respective strobe for AGP 8X.

To avoid additional signal mismatch, all of the lines within a group need to be the same type (either microstrip or stripline, but not both). This is because microstrip (surface traces) and striplines (buried traces) have different propagation velocities and mixing these can increase the flight time skew beyond acceptable limits. It is further recommended to route all signals within a group on the same layer. Routing studies have shown that these guidelines can be met. The maximum trace length requirements must not be violated by any signal. It is recommended to match the signals as closely as possible to provide timing margin.

1.6.5.3 Strobe Trace Routing Considerations

Since the strobe signals (**AD_STB0**, **AD_STBF0**, **AD_STB0#**, **AD_STBS0**, **AD_STB1**, **AD_STBF1**, **AD_STB1#**, **AD_STBS1**, **SB_STB**, **SB_STBF**, **SB_STB#**, and **AD_STBS**) act as clocks on the source synchronous AGP interface, special care should be taken when routing these signals. Note that the complement signals are available in AGP 4X mode only. (The AGP 8X strobes with “F” and “S” are not truly complement signals.) For AGP 2X, it is recommended that the strobe signals be routed in a 1:2 trace width/space ratio or greater relative to the other signals of the bus. For example, in a bus where the other signals are routed with 5 mil trace width and 5 mil spaces between, the strobe signals would be routed with a 10 mil separation on both sides. For AGP 4X, the strobe signals should be routed together (no other intervening signals) and separated by at least 4 times the maximum dielectric thickness (~5 mils for microstrip 4 layer board). The spacing from the strobe to an adjacent signal should also be 4 times the maximum dielectric thickness. In an AGP 8X design, the strobe signals should be routed together also but separated by at least five times the maximum dielectric thickness to each other and to any other signal routed adjacent to the strobe. This recommendation is intended to reduce the crosstalk noise coupled onto the strobes from other signals on the bus, as well as to reduce the noise coupled from the strobe signals onto adjacent lines. The strobe pair route lengths should be matched as closely as possible (< 10 mils). For AGP 8X the pair route lengths should be matched to < 5 mils delta.

1.6.5.4 Clock Routing and Skew

Clock skew between the AGP compliant graphics chip and the core chipset must be held to under 1 ns. The motherboard is allotted 0.9 ns of which a portion must be allotted to the clock generator circuit. The add-in card is allotted 0.1 ns of the total skew requirement. Figure 38 shows the trace segments that make up the clock routing.

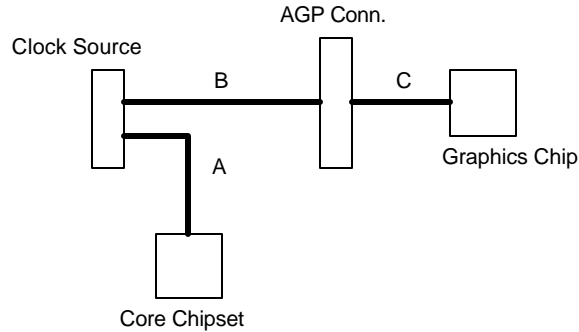


Figure 38: Required Trace Length to Minimize Clock Skew

The following clock skew equation is intended for use by the motherboard designer since the clock delay on the add-in card is tightly regulated to $0.6 \text{ ns} \pm 0.1 \text{ ns}$ by the AGP interface specification. It should be noted, however, that add-in card vendors should center the clock trace electrical length at 0.6 ns at typical board impedance so that variation in board impedance will be covered in the $\pm 0.1 \text{ ns}$ specification.

In this equation, all lengths are represented in terms of time (electrical length). This equation should be used several times to account for impedance variation on both the add-in card and the motherboard. The equation for clock skew is represented as a solution space as follows:

$$-(1.0 \text{ ns} - T_{\text{cgs}}) \leq A - (B + C + T_{\text{con}}) \leq (1.0 \text{ ns} - T_{\text{cgs}})$$

Where:

A = The electrical length of the clock trace on the motherboard from the clock generator to the Core Logic Memory Control (CLMC)

B = The electrical length of the clock trace on the motherboard from the clock generator to the AGP connector

C = The electrical length of the clock trace on the add-in card which from AGP interface specification is set at $0.6 \text{ ns} \pm 0.1 \text{ ns}$.

T_{con} = The delay through the AGP connector. We have found this to be 0.15 ns including crosstalk effects.

T_{cgs} = The clock generator skew between the clock signal to the chipset and the clock signal to the AGP connector.

NOTE: For a motherboard with graphics down, the above equation reduces to a simple subtraction of the two motherboard traces A and B.

The combination trace of $B + C + T_{\text{con}}$ should be matched as close to trace A as possible.

1.6.6 Interface Signaling Requirements

For AGP 2X, all signals are 3.3V compatible. 5V signals are not specified in the AGP signaling environment. All AGP 4X and 8X signals are +1.5V compatible. The master and target device must be capable of supporting the 3.3V signaling environment for AGP 2X and 1.5V signaling environment for AGP 4X and 8X. However, in a 1.5V signaling environment, there are some signals that are not 1.5V compatible. These signals are the interrupt lines, reset, and the AGP clock. The clock and reset signals are 3.3V level and are driven to both the target and master devices. The interrupt signals from the AGP bus must interface to the PCI bus interrupt controller. This controller and the PCI devices may be +5V devices. It is the requirement of the motherboard designer to properly interface the AGP interrupts to the PCI bus. Since the interrupt drivers are defined to be open-drain, this can be done by simply pulling up the PCI interrupts to 3.3V only, allowing the AGP interrupts to connect directly to the PCI interrupts.

Another option is to buffer the interrupts before they cross from the 3.3V domain to the 5V domain. Also, the two AGP interrupt lines must be “swizzled” with the PCI interrupt lines as described in the implementation note in the *PCI Local Bus Specification, Revision 2.1*. As an example, in a system with three PCI slots (and one AGP device or slot), the interrupts should be connected such that **INTA#** of each slot is assigned to a unique input on the system interrupt controller.

1.6.7 IDSEL Routing for Add-in Cards

As it is described in the *AGP Interface Specification, Revision 2.0*, **IDSEL** is not a pin on the AGP connector.

When the graphics device on the add-in card is designed for exclusive operation on the AGP interface, the device does not have an external **IDSEL** pin. There is no connection of an **IDSEL** signal to any **AD** signal on the add-in card. In this implementation, the device asserts **DEVSEL#** based monitoring of the **AGP** interface inside the graphics controller.

When the graphics device on the add-in card is to be used on both AGP and PCI bus segments, then the device needs to have two modes of operation. When in the AGP mode, it generates **DEVSEL#** as described in the AGP only implementation above. When used in a PCI mode of operation, the device must provide an external **IDSEL** that is connected to one of the **AD** signals in the system. This **IDSEL** signal is not connected externally to an **AD** signal on add-in cards designed for use on the AGP interface.

1.7 IC Packaging Considerations

This section describes some of the important signaling characteristics that packaging designers need to take into consideration before designing a package. These include conductor characteristics, conductor coupling, package layers, and die bonding method and inter-bond coupling.

Quad flat pack (QFP) packages are inexpensive, but the signaling leads usually have no reference plane. This causes the leads to be very inductive and to have significant crosstalk. The more signals there are the larger the package must be, which makes the leads longer and more inductive. The power leads are also inductive, increasing the impedance of the power delivery path. QFP is not a recommended package type for any AGP speed above 2X.

The ball grid array (BGA) package has a much higher signal density interconnect allowing much shorter, lower inductance internal lead lengths. Multi-layer BGA packages are recommended to allow all signaling paths (especially the high-speed source synchronous signals) to have a good ground reference plane. This gives the signal a good signal return path (assuming a good pinout). This also reduces signal lead inductance and makes the internal leads into transmission lines. The characteristic impedance of the transmission lines can be tuned to minimize signaling discontinuities between the package and the board. The transmission lines also have less crosstalk than leads with no reference plane. Finally, the multi-layer package allows better, lower inductance power delivery to the die. This improves simultaneous switching output (SSO) noise effects.

Most BGA packages use wirebonding for getting the signal from the package the die. Wirebonds have significant inductance of their own. Wirebonds are closely spaced, and there is also no close reference plane, so wirebonds have a significant amount of crosstalk. Their inductance hurts both signaling and power delivery. Reducing the crosstalk and improving power delivery requires many extra power and ground bond wires (see section 1.7.3). Wirebond BGA is suitable for speeds up to 4X. AGP8X can be supported in wirebond only if the special precautions (grounded wirebonds between each source synchronous signal) described below are taken.

A high-performance option for package technology is flip-chip FC-BGA. This package provides significant improvement to wire-bond BGA technology by decreasing the overall signal inductance and crosstalk. Short, wide bumps replace the long wires. Flip-chip packaging is recommended for AGP8X because of the improved crosstalk and power delivery, and increases the voltage and timing margins. While it is possible to use wirebond packaging for AGP8X, the cost of the additional ground and power wires can be high.

1.7.1 Signal Skew

Time skew between signals depends on differences between signal loading and internal lead length, and may also depend on signal transitions on neighboring traces. Signal transitions must charge and discharge the trace capacitance. Two signals driving different capacitive loads will have different timings. Neighboring signals switching in-phase or out-of-phase may also have a significant effect if coupling or crosstalk is large.

Mismatched package lead lengths contribute to skews because of their flight time differences. The timing budgets are based on a package with signal leads that are routed radially out from the die pad to the pin and have very similar flight times. AGP device vendors are encouraged to keep these lead lengths as short and as equal in length as possible. (In a BGA design, the inner and outer row of balls will be necessarily different in length by the depth of the ball array. This depth should account for most of the lead length difference.) Package vendors should publish any significant differences ($> 10\text{pS}$) in package lead flight times so that the board layout can compensate.

AGP 8X defines its timings pad-to-pad, therefore a package trace length document should accompany the core logic and the graphic controller chips. This document will help designers better match signal skewing at the motherboard and add-in card level.

1.7.2 Signal Loading

The AGP specification has limits on the package capacitance and recommendations for package inductance. It also requires that the source synchronous signals have package capacitive loading as closely matched as possible. This is a function of trace length, proximity to other metal structures, reference planes, and the relative dielectric strength. Package pin capacitance and induction are very design dependent.

In QFP, pin capacitance is a strong function of the lead length. This is due to the fact that lead spacing is generally constant among the leads. The range between minimum and maximum capacitance will depend on the overall size of the QFP.

In BGA packaging, the capacitance will depend on the existence of a power or ground plane, the width of the signal trace, the density of signal traces, the width of the plating bars, and the density of the plating bars. Traces routed over a good reference plane can be treated as transmission lines that do not contribute to a lumped capacitive or inductive component in the package. In general, signals routed to the corners of the pin field tend to be longer and be in areas of higher metal density. This makes the corner pins a poor choice for source synchronous signals which need close package route and capacitance matching. Also, BGA packages often have “plating bars” from the ball connection to the outside edge of the package to allow the outer layers to be plated after construction. These plating bars are generally short enough relative to the signal rise and fall time to look like parasitic capacitance on the ball. This has to be added into the pin capacitance mismatch.

1.7.3 Signal Routing and Crosstalk

Good signal integrity and timing depends on minimal crosstalk. Crosstalk between two traces in a package is no different than crosstalk on the board. Crosstalk between two traces is a function of the coupled length, the distance separating the traces, the “aggressor” signal edge rate, and the degree of mutual capacitance and inductance. In a non-homogeneous medium (microstrip traces), both forward and backward crosstalk will be excited on the victim signal trace, whereas only backward crosstalk will be excited within a homogeneous medium (stripline traces). Crosstalk is a problem for both asynchronous and synchronous signals. The result can be false triggering or data corruption. Mutual capacitance and inductance between signal traces can be reduced by the presence of a ground plane.

A ground plane is effective for traces in the package itself, but bondwires are too far from any reference plane to gain any benefit. The self-inductance loop area formed by a bondwire has a large cross section, and adjacent bondwires can couple very effectively, with bondwires 5 or 10 wires away contributing significant crosstalk. The crosstalk can be reduced if ground wires are introduced between the signals to act as return paths and shields. Bondwires that do not travel in the same plane (or arc) as the signal

wires will not be tightly coupled to them and cannot act as effective return paths and shields. Figure 39 and Figure 40 illustrate the use of shielding bondwires.




Figure 39: Interleaved Ground Bondwires for Signal Return and Shielding




Figure 40: Cross Section of Bondwires With Signal and Shield Ground Wires In Same Plane

1.7.4 Power Pins in the Signal Ball Array

Source synchronous signals have small timing and voltage margins and therefore need a good current return path to avoid impedance discontinuities and signal quality problems. This means that there must be good power return paths in the package. Power pins should be placed in the signal array to reduce the loop area that will cause the signal path to become overly inductive and limit signal propagation. The placement of power pins depends on package construction. In QFP, the best arrangement is to interleave signals with a power or ground signal. The strong mutual coupling of the adjacent leads will

reduce the self-inductance, much like the interleaved bondwire diagrams above.

In BGA packages with power planes, the plane over which the signals are routed should have frequent pins connected to the same plane on the board. Every signal pin should have an adjacent power pin (see Figure 41A and Figure 42). This minimizes the inductive loop area. It is recommended that the ground plane be used for the signal routing reference. Since the signal is referenced to only one signal plane, bypass capacitance is needed to allow return currents on the other plane to complete the loop. On-die capacitance is especially effective to aid signal return currents. If the package allows stripline routing, then the signals can be referenced to both Vddq and ground. In this case, both planes need to be frequently connected to their respective board planes (Figure 41B).




Figure 41: BGA Package Ballout Examples for AGP2X and AGP4X

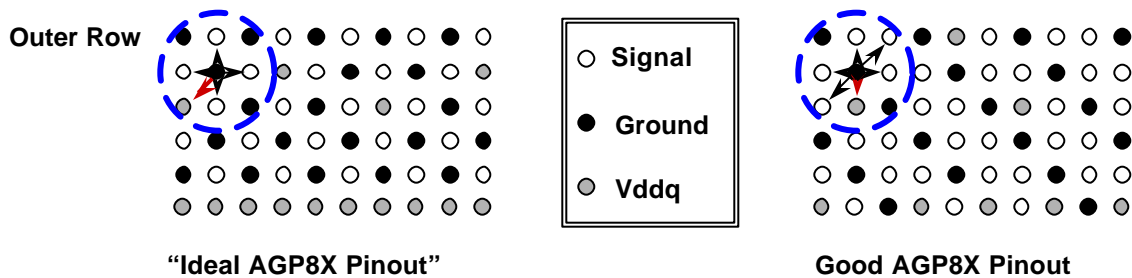


Figure 42: BGA Package Ballout Examples for AGP8X

2 AGP Signals Sensitivity Analysis and Measurement Techniques

These guidelines are primarily for AGP system designs that use an AGP-compliant graphics controller on an AGP-compliant add-in card. Designs with both AGP-compliant components mounted on the motherboard can also use these techniques. In this chapter, the terms 1X, 2X, 4X, and 8X modes refer to a specific bus electrical signaling and transfer speed according to the following table:

Table 16: AGP transfer mode definition for Design Guide

Mode	Transfer speed	Electrical signaling	Definition reference
1X mode	66MT/s	3.3V or 1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
2X mode	133MT/s	3.3V or 1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
4X mode	266MT/s	1.5V signal swing.	<i>AGP Interface Specification, Revision 2.0</i>
8X mode	533MT/s	0.8V signal swing.	<i>AGP3.0 Interface Specification, Revision 1.0</i>

This chapter is divided into three major sections:

- Section 2.1 deals with designs that are AGP 1X transfer mode (1X mode) and AGP 2X transfer mode (2X mode) specific. This includes both 3.3 volt signaling and 1.5 volt signaling.
- Section 2.2 describes AGP 4X transfer mode (4X mode). 4X mode designs need to also operate in 2X mode. 2X mode compliance will be achieved when a design is 4X mode compliant.
- Section 2.3 describes AGP 8X transfer mode (8X mode). This captures the most stringent layout guidelines in order to achieve 8X mode. The board layout design using the 8X guidelines will be 4X and 2X mode operations compliant.

2.1 2x Mode Board Design Guidelines

2.1.1 2X Mode Routing Rules

All of the techniques used in this section work for 1.5 volt signaling and 3.3 volt signaling. These board design guidelines were first based on 3.3 volt signaling. Analysis of 1.5 volt signaling did not result in reductions of any margins reported in this section.

These guidelines can be used three ways:

- The line length recommendations can be followed as they are described in Table 18 and Table 19. These rules allow enough variation to meet most designs.
- Section 2.1.3 can be used to apply variations to the existing design recommendations. This might be desired to take advantage of margin in other parts of a design (for example, lines routed shorter than those assumed by this guideline might not need to be matched as tightly). This can be done by applying the timing numbers shown in Table 26 or Table 27 to your own design.
- Section 2.1.2 can be used as a pattern for additional simulations, which may then be used for other designs. This section describes how the analysis was performed to make these guidelines. This can yield the most flexible designs; however, it requires extensive simulations.

2.1.1.1 AGP 2X Mode Design Considerations

With source synchronous data transfers (AGP 2X mode), skew between traces is a major interconnect design consideration. Skew is defined in the formula $T_{\text{skew}} = T_{\text{data}} - T_{\text{strobe}}$ where T_{data} and T_{strobe} are the flight times of the data and strobes respectively. Various effects that include line length mismatch, capacitance loading variations, crosstalk, and the voltage level from which a driver starts driving, cause these skews.

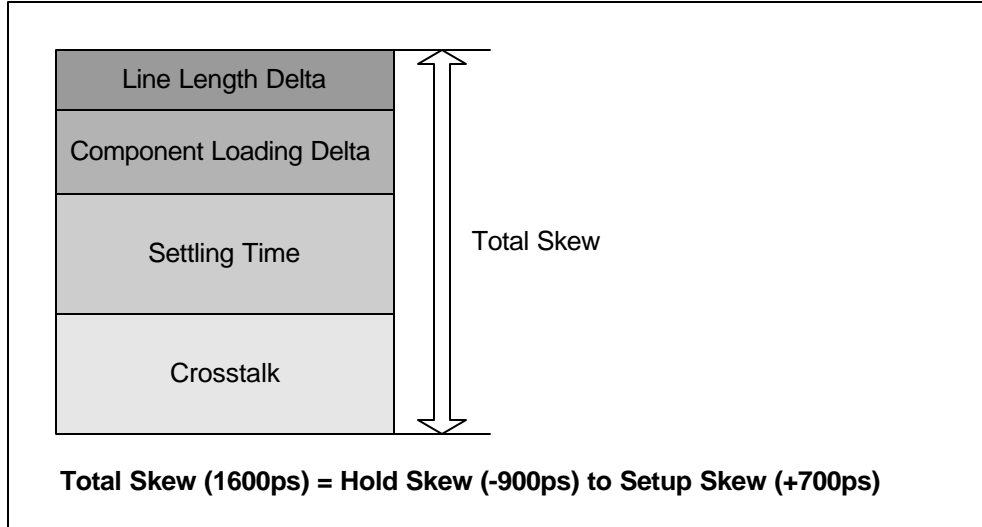


Figure 43: Example of Skews

Figure 43 is an example of how some of the different components contribute to the total skew. Line length mismatch is the physical difference in line lengths. Component loading mismatch is the difference in capacitive loading. These two effects are mainly independent of other effects.

Settling time and crosstalk have the largest effect on skews. Settling time creates uncertainty in the signal transition starts. Crosstalk is the effect of coupling between traces. Both of these interact with each other and with some of the other parameters in the interconnect.

With AGP 2X mode, it becomes more difficult to partition the timings between motherboard and add-in card. This is due to the fact that crosstalk and its effect on impedance discontinuity greatly increases the electrical skew between traces.

Throughout Section 2.1 of this document, the term “data” refers to the AD[31::0], C/BE[3::0]# and SAB[7::0] signals. The term “strobe” refers to AD_STB[1::0] and SB_STB. However, data timing and layout relationships will always be restricted to one of three groups of data as seen in Table 17. The strobe is restricted to its associated data group.

Table 17: Data and Associated Strobes

AGP 2X / 4X		AGP 8X	
<i>Strobes</i>	<i>Data Group Association</i>	<i>Strobes</i>	<i>Data Group Association</i>
AD_STB0* AD_STB0#*	AD[15::0] C/BE[1::0]#	AD_STBF0 AD_STBS0	AD[15::0] C#/BE[1::0]
AD_STB1* AD_STB1#*	AD[31::16] C/BE[3::2]#	AD_STBF1 AD_STBS1	AD[31::16] C#/BE[3::2] DBI_HI, DBI_LO
SB_STB* SB_STB#*	SBA[7::0]	SB_STBF SB_STBS	SBA#[7::0]

* The complement strobe signals are available in the AGP 4X and 8X modes only.

2.1.1.2 Source Synchronous Recommendations for Add-in Cards

Table 18: Add-in Card Recommendations

Width:Space	Zo	Trace	Line Length	Line Length Matching
1:2	50 Ω to 85 Ω	Data / Strobe	0.0 in < line length < 3.0 in	Strobe ± 0.5 in of group

All of the data line lengths within a group of signals need to be within ± 0.5 inches of their associated strobe. The board impedance needs to be in the range of 50 Ω to 85 Ω . This range is used to cover design targets and manufacturing tolerances.

Because crosstalk is a large component of skew, it is necessary to specify board routing. All traces need to be routed with a separation of two times the trace width. Additionally, all lines within a group need to be of the same type (either microstrip or stripline). This is because microstrip (surface traces) and striplines (buried traces) have different propagation velocities, and mixing these can increase the flight time skew beyond acceptable limits. Routing studies have shown that these guidelines can be met.

2.1.1.3 Source Synchronous Recommendations for Motherboards

The motherboard needs to have an impedance range of 50 Ω to 85 Ω . This range is used to cover design targets and manufacturing tolerances. All lines should be at least 1.0 inch long. The maximum line lengths are dependent on the type of trace and the amount of coupling.

The maximum line length is dependent on the routing rules used on the motherboard. These routing rules were created to give freedom for designs by making tradeoffs between signal coupling (trace spacing) and line lengths. These routing rules are divided by trace spacing. 1:1 spacing refers to the distance between the traces (air gap) as being the same width as the trace. 1:2 spacing refers to the distance between the traces as being twice the width of the trace.

For trace lengths that are less than 4.5 inches, a 1:1 data trace spacing gives a data line length mismatch of 0.5 inches with the strobe being the longest trace of the group. The strobe requires at least a 1:2 trace spacing. Use greater spacing, if possible, to minimize crosstalk and the possibility of coupling glitches from data to the strobe line. This is for designs that require less than 4.5 inches between the AGP connector and the AGP Target.

Longer lines have more crosstalk. Therefore, longer line lengths require a greater amount of spacing between traces to maintain skew timings. Table 19 shows 1:1 spacing may be used for lengths up to 4.5 inches, but 1:2 spacing allows lengths up to 9.5 inches. The line length mismatch is 0.5 inches with the strobe being the longest trace of the group. Here, the strobes are routed at 1:4, twice the spacing of the data lines. These timings also allow lines to neck down (reduced spacing) so that they may break out from a component.

Table 19: Motherboard Data / Strobe Line Length Recommendations

Width:Space	Zo	Trace	Line Length	Line Length Matching
1:1(Data) / 1:2 (Strobe)	50 Ω to 85 Ω	Data / Strobe	1.0 in < line length < 4.5 in	-0.5 in, strobe longest trace
1:2	50 Ω to 85 Ω	Data / Strobe	1.0 in < line length < 9.5 in	-0.5 in, strobe longest trace

In all cases, it is best to reduce the line length mismatch wherever possible to ensure added margin. It is also best to separate the traces, especially the strobe to neighboring traces, by as much distance as possible to reduce the amount of trace-to-trace coupling.

2.1.1.4 Control Signal and Clock Recommendations

Table 20: Control Signal Line Length Recommendations

Width:Space	Board	Trace	Line Length	Pull-up Stub Length
1:1	Motherboard	Control signals	1.0 in < line length < 8.5 in	< 0.5 in (Strobes < 0.1 in)
1:2	Motherboard	Control signals	1.0 in < line length < 10.0 in	< 0.5 in (Strobes < 0.1 in)
1:2 (1:4 to Strobe)	Motherboard	Clock		
1:2	Add-in Card	Control signals	0.0 in < line length < 3.0 in	
1:2 (1:4 to Strobe)	Add-in Card	Clock	4.0 in \pm 0.25 in (microstrip)	

Some of the control signals require pull-up resistors to be installed on the motherboard. The stub to these pull-up resistors needs to be controlled. The maximum stub length on a strobe trace is < 0.1 inch. The maximum stub trace length on all other traces is < 0.5 inches.

The clock lines on both the motherboard and the add-in card can couple with other traces. It is recommended that the clock spacing (air gap) be at least two times the trace width to any other traces. It is also strongly recommended that the clock spacing be at least four times the trace width to any strobes.

The clock lines on the motherboard need to be simulated to determine their proper line length. The motherboard needs to be designed to the type of clock driver that is being used and motherboard trace topology. These clocks need to meet the loading of the receiving device as well as the add-in trace length.

The add-in card trace length can be computed from the typical add-in card trace velocity.

2.1.2 Simulation Techniques

This section summarizes the interconnect simulations that were done to generate the routing guidelines for the AGP (66 MT/sec and 133 MT/sec) bus. The result of these simulations shows that designs can be done with buffers that meet the *AGP Interface Specification, Revision 2.0*. Figure 44 is a diagram of the topology used in this analysis.




Figure 44: AGP with One Connector Topology

The common clock simulations (66 MT/sec mode) limit the total line lengths (motherboard length plus add-in card length) to greater than 2.5 inches and less than 10.0 inches.

The source synchronous simulation results limit the total line lengths to 9.5 inches for the cases used. This line length is dependent on coupling between lines and line length mismatch.

The source synchronous simulation result limits the line length skews between a signal and its associated strobe on an add-in card ± 0.5 inch.

The source synchronous simulation results limit the motherboard's line length skews between a signal and its associated strobe to be 0.5 inches. The strobe needs to be the longest trace of each group.

The signal quality results for 3.3 volt signaling show overshoot and undershoot values approaching 2.1 V. Any device that connects to the AGP bus needs to be tolerant of these levels. Also, the ringback levels are well within the *AGP Interface Specification, Revision 2.0* requirements. Settling time was factored into both the common clock and the source synchronous results.

2.1.2.1 Methodology

Three main factors define the solution space: common clock flight times, source synchronous flight time skews, and signal quality. The common clock solution space is found by determining the topologies that result in lowest allowable flight times and the highest allowable flight times. The source synchronous solution space consists of all designs in which the flight time mismatch between a strobe and its associated data is less than the total allowable skew. Signal quality is checked to ensure that the solution space given by the common clock and source synchronous constraints meet all signal quality specifications.

Common clock (AGP 1X mode) flight times use worst case (extreme) buffers and interconnects to find its solution space. For minimum flight times, use a fast buffer driving the shortest (electrically) interconnects into a buffer with the smallest amount of loading. The physical line length is then increased until the minimum flight time is met. For maximum flight times, use a slow buffer driving the longest (electrically) interconnects into a buffer with the largest amount of loading. The physical line length is then varied until the maximum flight time is met.

Source synchronous (AGP 2X mode) flight time skews are found by comparing a strobe flight time with a data flight time for a particular interconnect, buffer, and loading. The skew is the total of both the setup skew and the hold skew. Setup skew is when the strobes' minimum flight time is compared to the data's maximum flight time at a given line length. Hold skew is when the strobes' maximum flight time is compared to the data's minimum flight time at a given line length. The results of these two comparisons must be less than the total allowable skew.

The results of the common clock solution space will give the minimum line lengths. The maximum line lengths will be determined by either the common clock or source synchronous. The signal quality of the signals within the solution space may constrain the solution space even further.

Of the three factors, source synchronous flight time skews are the most difficult to model. This is partly because they are a result of several parameters. These parameters include crosstalk, loading, impedance variations, line length variations, and settling time.

2.1.2.2 Buffer Models

The models for both the Master and the Target were identical. Both fast and slow buffer models were used in the simulations. In all cases, when a fast buffer was used, it was used in both locations. The same was also true for the slow buffers.

The buffer models were created from the *AGP Interface Specification, Revision 2.0*. These models also used power supply diodes. The package was modeled with an effective inductance of 10.0 nH.

The models also included an effective capacitance of 2 pF for the fast model and 8 pF for the slow model. When modeling the strobe, these capacitive values were varied by +2 pF and -1 pF.

2.1.2.3 Interconnect Models and Crosstalk Effects

Both the motherboard and the add-in card traces are modeled as transmission lines. The two parameters of importance in these models are the characteristic impedance (measured in ohms) and the propagation delay (measure in ns/ft, and sometimes loosely referred to as “velocity” even though this is not strictly correct).

The values used for characteristic impedance and propagation delay depend primarily on three contributing factors: trace geometry, printed circuit board manufacturing tolerance, and crosstalk from other traces. Trace geometries that are typically used for motherboards and add-in cards will yield impedances in the range of 60-75 Ω . Manufacturing tolerances can be expected to add another 15%, increasing this range to approximately 50-85 Ω .

The impact of crosstalk is a little more difficult to comprehend. To model the impact of crosstalk, it is necessary to understand how the impedance and velocity vary between the even mode (more correctly called the “common mode,” meaning that all coupled traces are switching in the same direction) and the odd mode (more correctly called the “differential mode,” meaning that coupled traces are switching in the opposite direction).

Common and differential mode impedances and propagation delays were modeled using a wide range of printed circuit board trace geometries (trace width, separation distance between traces, and trace height above ground/power plane). For this analysis, this crosstalk effect was included in the data traces. Strobe traces, however, did not include this effect because the strobes do not switch at the same time as the data traces. The results are summarized in Table 21.

After completing the simulations described in the next section, it was determined that the data traces must be designed so that their impedance is confined to 30-120 Ω (including all of the above effects). The geometries listed in Table 18, Table 19, and Table 20 achieve this goal.

The connector was modeled as a transmission line with 150 ps delay and 35-80 Ω characteristic impedance.

Table 21: Interconnect Impedance and Velocities

	Board Condition	Strobe		Data		Mode
		Z0 (Ω)	S0 (ns/ft)	Z0 (Ω)	S0 (ns/ft)	
6 mil x 12 mil Board	Fast	85	1.6	60	1.38	Odd
				99	1.70	Even
	Slow	50	2.0	45	1.72	Odd
				74	2.12	Even
6 mil x 6 mil Board	Fast	85	1.6	41	1.28	Odd
				119	1.80	Even
	Slow	50	2.0	30	1.60	Odd
				88	2.24	Even
Connector		80	2.0	63	1.83	
		35	2.0	52	2.17	

2.1.2.4 Simulation Setup

Simulations were run to find flight times for AGP 1X clocking mode, flight time skews for AGP 2X clocking mode, and signal quality for both modes. Earlier simulations had shown that the solution space would be confined by AGP 1X mode for the minimum line lengths and AGP 2X mode for the maximum line lengths.

All simulations were run using various combinations of line lengths for both the motherboard and the add-in card. The motherboard line lengths were varied from 0 to 10.0 inches in 0.5 inch increments. The add-in card line lengths were at 2.5 inches and 3.0 inches.

The strobe signal was simulated with eight different combinations of interconnects. These interconnects are fast/slow motherboard, fast/slow add-in card, and fast/slow connector.

The other signals were simulated with 32 different combinations of interconnects. These combinations are four velocity and impedance combinations of motherboard, four velocity and impedance combinations of add-in card and two combinations of connectors. The four combinations of board interconnect were fast and slow in both odd and even mode. The two combinations of connectors were fast odd mode and slow even mode. All of these simulations were done with both 1:1 and 1:2 motherboard trace spacing.

3.3 volt signaling simulations used 3.3 V for Voh and 0.0 V for Vol. A threshold value of 1.32 V was also used for all simulations. This 1.32 V was derived by taking 40% of the Voh minus Vol.

1.5 volt signaling simulations used 1.5 V for Voh and 0.0 V for Vol. A threshold value of 0.75 V (0.5Vddq) was also used for all simulations.

Additional simulations were run to find the clock to data out (T_{CO}) for the buffer models that were used. Section 2.1.2.6 describes how the T_{CO} simulations were used in the flight time and flight time skew calculations.

2.1.2.5 Simulation Parameters

In order to give designs the greatest amount of freedom, a number of assumptions were made. These assumptions were based on previous simulations, past experience, routing studies, and feedback from some vendors.

For the add-in card, routing studies have shown that a board can be routed with the following requirements:

- The maximum trace length of 3.0 inches.

- A line length skew of the data lines to be ± 0.5 inches with respect to their associated strobe.
- A trace spacing of 1:2, that is a 6 mil trace with 12 mil spacing (air gap) or a 5 mil trace with 10 mil spacing.

An additional assumption for the add-in card is that the board can be with an impedance range of 50 Ω to 85 Ω . These are thought to include the values of the most popular ranges of specified printed circuit boards (PCB).

For the motherboard, routing studies have shown that several different routing requirements are needed. For some designs, the trace spacing needs to be 1:1. Other designs need longer line lengths than can be achieved with a 1:1 spacing. These designs need to be made with wider spacing except where necessary to escape a component.

To allow for these routing needs, two motherboards each with a different trace spacing were used. These were with a motherboard spacing of 1:1 and 1:2. These motherboards were analyzed to determine the maximum line length that can be achieved for these spacings. All three of these styles of motherboards need to allow at least a 0.5 inches variation between the data lines. Routing studies have shown that it is usually best to have the strobe longer than the data.

An additional assumption for the motherboard is that the board can be specified with an impedance range of 50 Ω to 85 Ω .

Also all motherboards will be able to run in both AGP 1X mode and AGP 2X mode.

2.1.2.6 Flight Time

A previous analysis had shown that the minimum flight time occurs when the add-in card data line length would be at its shortest length (0.0 inches). Also, the minimum flight time is achieved with the fast buffer. Simulations were run with the various combinations of interconnect for the data lines on the motherboard and connector while varying the motherboard line length.

A previous analysis had shown that the maximum flight time occurs when the add-in card data line length would be at its longest length (3.0 inches). Also, the maximum flight time is achieved with the slow buffer.

By using the simulations that were required for the AGP 2X mode, it would be possible to find out what the maximum flight time is for a line length up to 10.0 inches. By using these results, the maximum line length on the motherboard for AGP 1X could be found.

Flight Time Measurement Methodology Chart




Figure 45: Flight Time Measurement

Figure 45 shows how flight time is measured with 3.3 volt signaling. In this example $T_{FLIGHT} = T_{VAL} - T_{CO}$ where T_{VAL} is the time that the signal crosses the threshold, and T_{CO} is the time that a test load crosses the threshold. The same technique is used with 1.5 volt signaling.

2.1.2.7 Flight Time Skew

The flight time skew simulations need to simulate any parameter that could cause a skew between two signals. The flight time skew simulations included varying the motherboard and add-in card line lengths, the effective capacitance in the buffer models, and crosstalk on each of the different interconnect combinations. They also included the effect that settling time (data pattern dependence) has on the flight time skew by comparing settled signals to unsettled signals.

For the source synchronous signals, the various parameters have a cumulative effect on the total flight time skew but do not add linearly. The maximum skew does not occur where each effect is maximized, and the manner in which each effect adds or interacts is not easily predicted. This means that a large number of cross variations need to be looked at.

In all cases, the strobe simulations were compared to the data simulations. Also, all flight time skew simulations were run with the add-in card line length at both 2.5 inches and 3.0 inches. This was to simulate the worst case variations between data and strobe attributed to the add-in card. The motherboard line length was varied from 0.0 inches to 10.0 inches in 0.5 inch increments. This was to find the maximum line lengths as they relate to skew.

All of the strobe simulations were done using the different types of interconnect combinations for motherboard, connector, and add-in card. These values did not include the effect of crosstalk. These simulations are grouped by motherboard routing rules. The values used with these rules are shown in Table 21. The add-in card and the motherboard use 1:1 and 1:2 spacing.

All of the data simulations were done using the different types of interconnect combinations for motherboard, connector, and add-in card. These values included the effect of crosstalk in the data simulations. These simulations are grouped by motherboard routing rules. The values used with these

rules are shown in Table 21. The add-in card used the 1:2 spacing, and the motherboard used 1:1 and 1:2 spacing.

The capacitive loading for the data simulations was held at the nominal values. The fast buffer used 2 pF and the slow buffer used 8 pF.

The capacitive loading was varied for the strobe simulations. These were varied -1 pF and +2 pF of the nominal values for both the fast and the slow buffers. This fast buffer had a nominal loading of 2 pF, so the strobe was simulated with an effective capacitance of 1 pF and 4 pF when the fast buffer was used. The slow buffer had a nominal loading of 8 pF, so the strobe was simulated with an effective capacitance of 7 pF and 10 pF when the slow buffer was used.

The strobe simulations were run only at a 7 ns cycle time. This is due to the fact that the strobe will constantly be switching states when transferring data. The number of cycles used was 4.

The data simulations were run varying the number of cycles depending on the cycle time. When the duration (half-cycle) was set to 7 ns, the number of cycles was set to 4. When the duration was set to 50 ns, the number of cycles was set to 1. These simulations were used to calculate the minimum and maximum flight times for signals that are settled and for signals that are running at the maximum transitional speed. This found any effects attributable to data pattern dependency.

In addition to finding the maximum accumulative flight time skew, the effect of each of the parameters needed to be analyzed. This was to find out what impact the different variables had on the total flight time skews.

2.1.2.8 Signal Quality

Settling time was not measured in these simulations but was taken into account with the flight time skew measurements.

Overshoot is measured at the highest voltage, and undershoot is measured at the lowest voltage at the receiver. Ringback from the rising edge is the lowest voltage that a receiver comes down to after crossing V_{ddq}. Ringback from the falling edge is the highest voltage that a receiver comes up to after crossing V_{ss}. Figure 46 is an example that shows how overshoot, undershoot, and ringback are measured with 3.3 volt signaling. The same method is applied to 1.5 volt signaling.

Signal Quality Measurement Methodology Chart





Figure 46: Signal Quality Measurements

2.1.2.9 Summary of Simulations Required

A sensitivity analysis that was done early in the design phase contributed to the reduction of the number of simulations required at this stage in the analysis. This analysis used 3-D graphs with two variables along the X-axis and Y-axis. The results of the simulations are displayed in the Z-axis. Also, pull down menus were used so that a variety of parameters could be viewed. As an example, one of these 3-D graphs is shown in Figure 47.



The number of simulations can be reduced even more by joining the simulations that are required for each of the three factors that define the solution space. The main contributor to the simulations required is the number of source synchronous simulations. These are the simulations that define the maximum line lengths for AGP 2X mode. For the common clock flight time simulations, an additional set of simulations is required to determine the minimum line lengths. All of the signal quality measurements can be taken from the source synchronous simulations. Table 22 is a table of these simulations. Each of these simulations needs to be done with all the variations of interconnects shown in Table 21.

Table 22: Summary of Simulations

Type	Interconnects	Driver	Loading	Duration	AC line lengths	MB line lengths
Data	1:2 spacing	Fast & slow buffers	Typical	7 ns & 50 ns	2.5 in & 3.0 in	0.0 in to 10.0 in
Data	1:1 spacing	Fast & slow buffers	Typical	7 ns & 50 ns	2.5 in & 3.0 in	0.0 in to 10.0 in
Data	1:1 spacing	Fast buffers	Typical	7 ns & 50 ns	0.0 in	0.0 in to 10.0 in
Strobe	typical	Fast & slow buffers	+2 pF & -1 pF	7 ns	2.5 in & 3.0 in	0.0 in to 10.0 in

2.1.3 Simulation Results

2.1.3.1 Flight Time

The results from the simulations were put into a spreadsheet for analysis. Data simulations were used for both the maximum and minimum flight times.

For the minimum flight time, the simulations used were those that had an add-in card line length of 0.0 inches. These used fast buffers and motherboard interconnects that were based on 1:1 spacing.

A spreadsheet was then used to calculate the flight times for each of these cases by subtracting the T_{co} from each flight time. The minimum values from these were found by taking the minimum value for a given motherboard line length and putting it into Table 23. These results show that the motherboard line lengths need to be greater than 1.0 inch.

Table 23: Minimum Flight Time

MB data line length (in)	0.0	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0
Min. data flight time (ps)	-99	-50	20	90	140	190	230	300	350	410	490	530	590	660	690	760	820	890	940

For the maximum flight time, the simulations used add-in card line lengths of 3.0 inches. These used slow buffers and motherboard interconnect spacing of 1:2.

A spreadsheet was then used to calculate the flight times for each of these cases by subtracting the T_{co} from each flight time. The maximum values found for a given motherboard line length were put into Table 24. These results show that the motherboard line lengths need to be less than 10.0 inches.

Table 24: Maximum Flight Time

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0
Max. data flight time (ns)	0.82	0.91	1.00	1.10	1.20	1.30	1.40	1.49	1.59	1.68	1.77	1.86	1.95	2.05	2.14	2.23	2.32	2.41	2.50

2.1.3.2 Flight Time Skew

The results from the source synchronous simulations were put into a spreadsheet for analysis. This spreadsheet compares the flight times from the strobe simulations to the data simulations.

Each case from all of the different strobe simulations was compared to the cases from several different data simulations. These cases are the minimum and maximum flight times for the rising and falling edges for each simulation. A summary of the comparisons that were done for each line length on the motherboard for a given line spacing is shown in Table 25.

Table 25: Summary of Source Synchronous Comparisons

Variable	Strobe	Data	Comments
Buffer location	Master / Target	Master / Target	compares use same location
Buffer type	fast / slow	fast / slow	compares use same type
Loading	-1 pF / + 2 pF	typical	compares all cases (2)
Cycle time	7 ns	7 ns / 50 ns	compares all cases (2)
Corner	rise / fall	rise / fall	compares all cases (4)
Flight time	min / max	min / max	compares all cases (4)
Fast motherboard	fast	fast odd / fast even	compares all cases (2)
Slow motherboard	slow	slow odd / slow even	compares all cases (2)
Connector	fast / slow	fast odd / slow even	compares all cases (2)
Fast add-in card	fast	fast odd / fast even	compares all cases (2)
Slow add-in card	slow	slow odd / slow even	compares all cases (2)
Add-in card line length	2.5 in	3.0 in	
Add-in card line length	3.0 in	2.5 in	

The results of all of these comparisons can be seen in Table 26 and Table 27.

Table 26 is a table of the maximum and minimum skews that resulted from a number of simulations with a motherboard spacing of 1:1. The top row is the data line lengths on the motherboard in inches. The next two rows are the setup skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches less than the data lines; the other is when they are of equal length. The last two rows are the hold skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches longer than the data lines, the other is when they are of equal length. The highlighted area is where the recommended trace lengths for motherboards were derived. The solution needs to be between -900 ps (hold skew limit) and 700 ps (setup skew limit).

Table 26: Skew From the Motherboard with 1:1 Trace Spacing

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5
Setup (ps) strobe -0.5 in	510	550	590	620	670	710	740	780	820	860	900	940
Setup (ps) strobe equal	410	460	510	550	600	630	570	700	750	790	820	860
Hold (ps) strobe equal	-470	-530	-570	-630	-690	-720	-770	-800	-820	-810	-850	-880
Hold (ps) strobe +0.5 in	-560	-610	-660	-730	-760	-800	-850	-860	-880	-900	-940	-960

Table 27 is a table of the maximum and minimum skews that resulted from a number of simulations with a motherboard spacing of 1:2. The top row is the data line lengths on the motherboard in inches. The next two rows are the setup skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches less than the data lines; the other is when they are of equal length. The last two rows are the hold skews in picoseconds between the strobe and the data lines. One of these is when the strobe is 0.5 inches longer than the data lines; the other is when they are of equal length. The highlighted area is where the recommended trace lengths for motherboards came from. The solution needs to be between -900 ps (hold skew limit) and 700 ps (setup skew limit).

Table 27: Skew from the Motherboard with 1:2 Trace Spacing

MB data line length (in)	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0
Setup (ps) strobe -0.5in	450	480	500	550	600	600	620	610	590	610	590	590	610	640	670	690	710	750	780
Setup (ps) strobe equal	360	390	440	490	530	530	520	520	510	520	500	520	550	590	620	640	660	690	720
Hold (ps) strobe equal	-470	-510	-540	-650	-710	-730	-760	-730	-740	-700	-700	-710	-720	-740	-760	-770	-800	-820	-840
Hold (ps) strobe +0.5in	-550	-590	-670	-750	-780	-810	-810	-810	-780	-770	-790	-790	-810	-820	-840	-860	-880	-900	

2.1.3.3 Signal Quality

The analysis of signal quality was done by finding the minimum and maximum values from the different parameters with fast buffers. The greatest amount of overshoot was 5.25 V, and the greatest amount of undershoot was -2.05 V. The worst case ringback from the rising edge was 2.62 V, and the worst case ringback from the falling edge was 0.90 V. These values are for 3.3 volt signaling.

2.2 4X Mode Board Design Guidelines

2.2.1 4X Mode Routing Rules

These guidelines are primarily for AGP 4X transfer mode (4X mode) designs that use a 4X mode compliant graphics controller on a 4X mode compliant add-in card that plugs into a 4X capable motherboard with a 4X compliant system controller. Designs with both 4X compliant components on the motherboard can also use these techniques.

Add-in card designs must follow length rules that are specified in the *AGP Interface Specification, Revision 2.0*. These rules are reiterated in Table 30.

For motherboard designs, there are three ways that these guidelines can be used:

- Designers who do not have the time or resources to complete a full simulation of their design can route their layout following the line length rules in Table 29, Table 30, and Table 31. These rules preserve enough timing margin to ensure a working design without requiring further simulations. The rules are flexible enough to be used on most designs. Care must be taken to follow these rules precisely, if no further simulations are to be used.
- Designers wishing to create their own set of line length rules can use the information included in Section 2.2.3. This might be desired in order to take advantage of margin in other parts of a design (for example, lines routed shorter than those assumed by this guideline might not need to be matched as tightly). This can be done by reviewing the results that are shown in Figure 54, Figure 57, Figure 58, Figure 59, and Figure 60 and applying them to your own design.
- Designers who wish to base their designs on their own simulations can follow Section 2.2.2 as a pattern for additional simulations, which may then be used for other designs. This section describes how the analysis was done to make these guidelines. This will yield the most flexible designs; however, it requires extensive simulations.

Sections 2.2.4 and 2.2.5 give additional design information that can be used with all three methods of designing a 4X mode compliant motherboard and add-in card.

2.2.1.1 AGP 4X Mode Design Consideration

Scaling the AGP interconnect from 2X transfer mode (2X mode) to 4X mode presents some interesting challenges. These challenges include a two-fold increase in bus speed and reducing the amount of interconnect skew by 1/3.

Reductions are required in several of the parameters that make up the total skew budget. An example of this is illustrated in Figure 48. Line lengths require closer matching in 4X mode than in 2X mode. The loading variation must be reduced in 4X mode. The crosstalk contribution to skew must be reduced by

the greatest amount in order to avoid excess skew due to increased inter-symbol interference (ISI). This means that the separation distance between 2X and 4X printed circuit board traces is increased.

Though not shown in Figure 48, Inter-Symbol Interference (ISI) is much worse in 4X mode when using the same interconnect parameters as were used in 2X mode, because the shorter cycle time makes these signals more sensitive to oscillations and ringing.

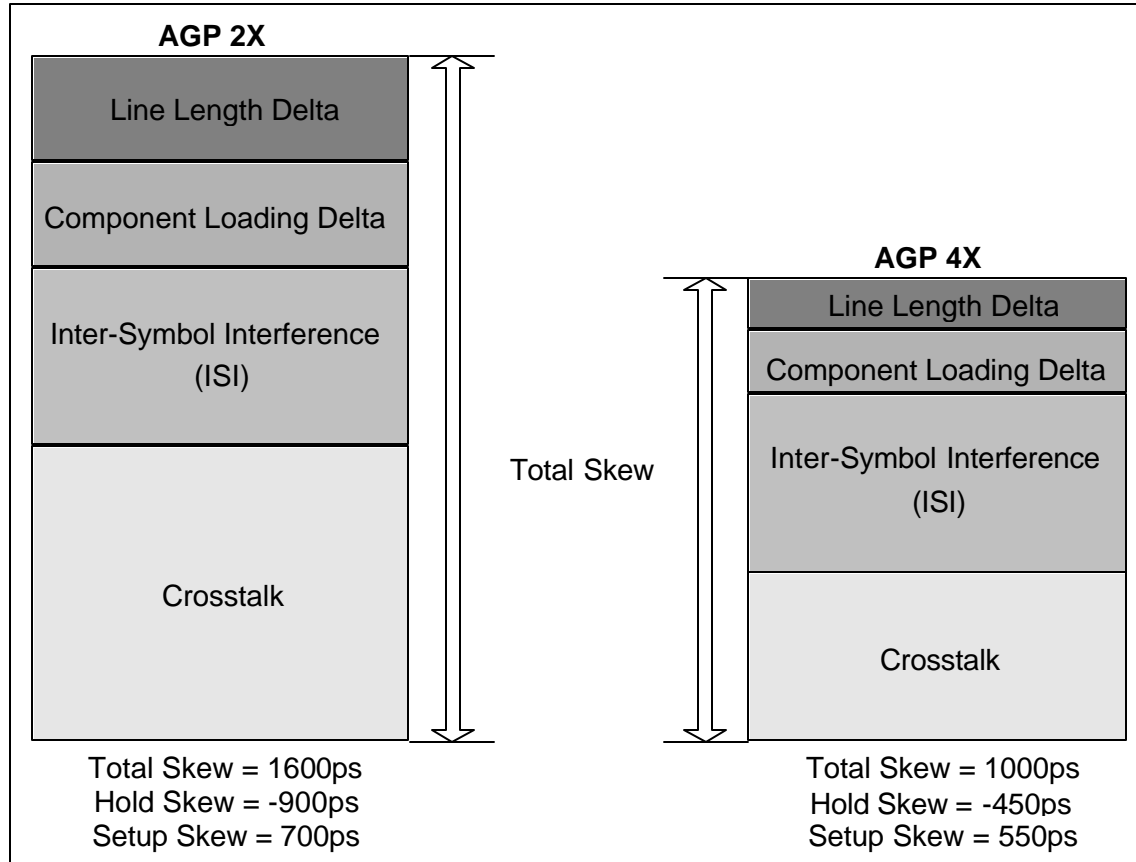


Figure 48: Example of Skew Differences Between 2X Mode and 4X Mode

Throughout Section 2.1 of this document, the term “data” refers to the **AD[31::0]**, **C/BE[3::0]#** and **SAB[7::0]** signals. The term “strobe” refers to **AD_STB[1::0]** and **SB_STB**. However, data timing and layout relationships will always be restricted to one of three groups of data as seen in Table 28. The strobe is restricted to its associated data group.

Table 28: 4X Mode Data and Associated Strobe / Strobe#

Data	Associated Strobe	Associated Strobe#
AD[15::0] and C/BE[1::0]#	AD_STB0	AD_STB0#
AD[31::16] and C/BE[3::2]#	AD_STB1	AD_STB1#
SBA[7::0]	SB_STB	SB_STB#

2.2.1.2 Source Synchronous Rules for Motherboards

Table 26 gives the routing rules for motherboards. Boards laid out following these rules will work robustly, requiring minimum additional simulations.

Table 29: Motherboard Recommendations

Width: Space Minimum ^{1, 2}	Z0	Trace	Line Length ³	Line Length Matching
1:4	54 Ω to 66 Ω	Strobe to Strobe#	1.0 in \leq line length \leq 7.25 in	< 0.1 in
1:6	54 Ω to 66 Ω	Strobe to Data	1.0 in \leq line length \leq 7.25 in	(see below)
1:4	54 Ω to 66 Ω	Data	1.0 in \leq line length \leq 7.25 in	Strobe and Strobe# \pm 0.125 in of group
1:3	54 Ω to 66 Ω	Data	1.0 in \leq line length \leq 6.0 in	Strobe and Strobe# \pm 0.5 in of group

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups, scale the trace width with the dielectric thickness. For more information, refer to Section 2.2.1.6.
3. Line lengths are based on microstrips (external traces). Striplines (internal traces) will require a 25% reduction in line lengths. This is due to a 25% reduction in propagation velocity.

The recommended motherboard impedance range is 60 $\Omega \pm 10\%$. This range plus the width:space ratios in Table 29 allow sufficient resistance to crosstalk effects. The stackup shown in Section 2.2.4.1 makes the required coupled impedance range of 48 Ω to 73 Ω possible. This range is used to cover design targets and manufacturing tolerances. All lines should be at least 1.0 inch in length to ensure conformance to the AGP 2X design guidelines. The maximum line lengths are dependent on the type of trace and the amount of coupling.

The maximum line length is dependent on the routing rules used on the motherboard. These routing rules are intended to give freedom for designs by making tradeoffs between signal coupling (trace spacing) and line lengths. The routing rules are divided by trace spacing. 1:3 spacing refers to the distance between the traces (air gap) as being three times the width of the trace. 1:4 spacing refers to the distance between the traces as being four times the width of the trace.

In all cases, it is best to reduce the line length mismatch wherever possible to ensure added margin. It is also best to separate the traces by as much as possible to reduce the amount of trace-to-trace coupling.

2.2.1.3 Source Synchronous Rules for Add-in Cards

Table 30 gives the routing rules for add-in cards.

Table 30: Add-in Card Rules

Width:Space Minimum ^{1,2}	Z ₀	Trace	Line Length ³	Line Length Matching
1:3	54 Ω to 66 Ω	Strobe to Strobe#		< 0.1 in
1:6	54 Ω to 66 Ω	Strobe to Data		
1:3	54 Ω to 66 Ω	Data	≤ 1.5 in	Strobe and Strobe# ± 0.25 in of group

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups, scale the trace width with the dielectric thickness. For more information, refer to Section 2.2.1.6.
3. Line lengths are based on microstrips (external traces). Striplines (internal traces) will require a 25% reduction in line lengths. This is due to a 25% reduction in propagation velocity.

The required add-in card impedance range is $60\ \Omega \pm 10\%$, therefore, it will be necessary to use a controlled impedance board. This range is used to cover design targets and manufacturing tolerances.

The routing rules are divided by trace spacing. 1:3 spacing refers to the distance between the traces (air gap) as being three times the width of the trace. 1:4 spacing refers to the distance between the traces as being four times the width of the trace.

In all cases, it is best to reduce the line length mismatch wherever possible to ensure added margin. It is also best to separate the traces by as much as possible to reduce the amount of trace to trace coupling.

2.2.1.4 Control Signal and Clock Recommendations

Table 28 gives the routing rules for the control signals and clock lines.

Table 31: Control Signal Line Length Recommendations

Width:Space Minimum ^{1,2}	Board	Trace	Line Length	Pull-up / Pull-down Stub Length
1:3	Motherboard	Control Signals	$0.0 < \text{line length} < 7.5\text{in}$	< 0.1in
1:4	Motherboard	Clock		
1:3	Add-in Card	Control Signals	$0.0 < \text{line length} < 3.0\text{ in}$	
1:4	Add-in Card	Clock	$4.0\text{ in} \pm 0.25\text{ in}$	

Notes:

1. These ratios assume a particular board stackup. For more information, refer to Section 2.2.4.1.
2. For different stackups scale the trace width with the dielectric thickness. For more information, refer to Section 2.2.1.6.

Some of the control signals require pull-up or pull-down resistors. These may be placed anywhere along the trace, but must be routed so they do not introduce any stub greater than 0.1 inch. To meet this requirement, it is strongly recommended that discrete pull-ups be used.

The clock lines on the motherboard should be simulated to determine their proper line length. The motherboard must be designed to the type of clock driver that is being used. In all cases, it is recommended to have a trace spacing that is at least 4 times the trace width.

2.2.1.5 Inter-Symbol Interference

Inter-Symbol Interference (ISI) affects signals when the driver does not reach steady state prior to the next switching cycle. This causes the flight times to vary from one transition to another. This variability will increase the amount of skew between the strobe and data signals.

Figure 49 is an example of ISI. This case shows a skew of 150 ps due to a data signal transitioning with a bit time of 3.75 ns and another at 11.25 ns (3 bit times). This example shows how flight time can change depending on bit pattern. The narrow pulse (representing a 1010... bit pattern) has a delay which is noticeably different from the wide pulse (representing a 111000... bit pattern). This difference of 150 ps must be included as skew.




Figure 49: Example of ISI

ISI is caused by signal ringing; therefore, some methods used to reduce the amount of ISI are active receiver clamps, termination, and driver impedance matching. Clamps are difficult to implement because of the low signaling levels used. There is insufficient voltage to develop good bias level for "early clamps," and other types of feedback clamps affect the flight time and may introduce too much skew.

Terminations add cost to the printed circuit boards (PCB) and parts. The cost added to the boards would be due to routing that would in all likelihood increase the layer count. They are also not backward compatible to the AGP 2X signaling scheme.

Driver impedance matching works quite well if the buffer impedance is matched exactly to the interconnect impedance. This results in the signal at the receiver going to the rail without any overshoot or reflections. This ideal circuit is impossible to implement because of manufacturing tolerances and impedance changes in the interconnect due to crosstalk. However, reductions in the range that impedance varies in the buffer and the interconnect are possible and can be used to reduce timing skew.

The buffer's slew rate can also be a large factor of ISI. This happens when the signal's round trip propagation delay plus the total rise time (slew rate) is greater than the period. This is due to the interference of the returning wave.

The following sections describe interconnects and buffer design techniques for reducing ISI and other timing skew sources.

2.2.1.6 Board Constraints

The two main factors that contribute to PCB impedance variations are the manufacturing tolerances and crosstalk. The range over which these factors vary must be reduced to decrease the total impedance variation.

Reducing the board tolerances is achieved by specifying the board and layer stackup to board vendors. For AGP 2X mode, this range was 50 Ω to 85 Ω . For 4X mode, the range is reduced to 60 $\Omega \pm 10\%$.

Reducing the amount of coupling between driven lines will reduce the impedance variation due to crosstalk. Sections 2.2.2.4 and 2.2.2.5 describe what these values are and how they were derived. In a nutshell, crosstalk can be reduced by decreasing the separation distance of a trace to the nearest reference plane (dielectric thickness) and/or increasing the separation distance to adjacent traces (trace spacing).

Throughout Section 2.2, the spacing between traces is specified as a ratio of the trace width (width:space) which is a commonly used method to specify a board layout. This ratio is correct for the recommended board stackup, but the desired result is a ratio of trace spacing to dielectric thickness. Appendix A contains a more in depth explanation of crosstalk and how to include it in simulations.

All 4X mode signals, along with the clocks and strobes, are not to cross any plane splits. (These might occur if a plane is used to supply two separate Vcc voltages.) If the common clock signals cross split planes, they should be simulated to insure that signal quality is not compromised. For more information on the effect that split planes have on the traces that cross them, refer to a paper *Signal Integrity Issues at Split Ground and Power Planes* presented by Haw-Jyh Liaw and Henri Merkelo at the 46th Electronic Components & Technology Conference 1996 (IEEE Catalog No. 96CH35931).

2.2.2 Simulation Techniques

This section summarizes the interconnect simulations that were done to generate the routing guidelines for the AGP (266 MT/s) bus. The result of these simulations shows that designs using buffers defined in the *AGP Interface Specification, Revision 2.0* can be done. Figure 50 is a diagram of the topology used in this analysis.

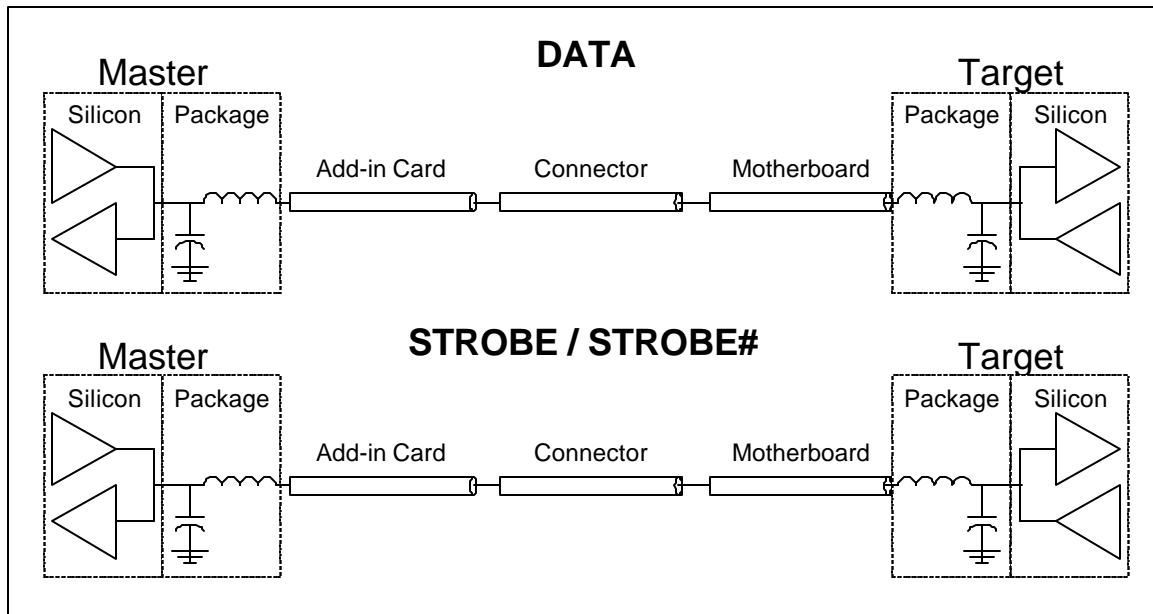


Figure 50: AGP Source Synchronous Topology

The source synchronous simulation results limit the total line length to 7.25 inches for the simulated cases. Refer to Section 2.2.3.4 for a description of these cases. This line length is dependent on coupling between lines and line length mismatch.

Source synchronous signaling requires limits on the amount of variation between the strobe/strobe# line length and the data lines for a signal group. This variation is recommended to be no more than ± 0.125 inch from any data within the group to either of its strobes (strobe, strobe#).

The signal quality results show overshoot and undershoot values approaching 0.7 volts. Also, the ringback levels are well within the *AGP Interface Specification, Revision 2.0* requirements.

2.2.2.1 Methodology

Three main factors define the range of trace lengths useable for AGP 4X: common clock flight times, source synchronous flight time skews, and signal quality. As shown in Figure 51, common clock flight times limit the length to a minimum of 1.0 inches. The maximum line lengths are restricted by flight time skews.




Figure 51: Solution Space Plot

Figure 53 is actually a “solution space” plot formed by combining all of the simulation results into one big “pass/fail” test. This is a pass/fail plot of the simulation results compared to the limits (flight time, flight time skew and signal quality) that are included in the *AGP Interface Specification, Revision, 2.0*. Some of the data that made up this plot is shown in Appendix A as 3D plots. Appendix A also includes data that was not used for this solution.

Common clock (AGP 1X mode) flight times use worst case (extreme) buffers and interconnects to determine the solution space. A fast buffer driving the shortest (electrically) interconnects into a buffer with the smallest amount of loading produces minimum flight times. The physical line length is varied until the minimum flight time is met. A slow buffer driving the electrically longest interconnects into a buffer with the largest amount of loading produces maximum flight times. The physical line length is then varied until the maximum flight time is met.

Source synchronous (AGP 2X mode and AGP 4X mode) flight time skews are found by comparing a strobe flight time with a data flight time for a particular interconnect, buffer, and loading. Setup skew occurs when the strobe’s minimum flight time is compared to the data maximum flight time at a given line length. Hold skew occurs when the strobe maximum flight time is compared to the data minimum flight time at a given line length. The results of these two comparisons must be less than the total allowable skew.

Signals involved in source synchronous transfers must meet line length restrictions set by minimum and maximum flight times as well as flight time skew. The other “outer loop” common clock signals are limited only by the minimum and maximum flight times. The signal quality of the signals within the solution space may constrain the solution space even further.

Of the three factors, source synchronous flight time skews are the most difficult to model. This is partly because they are a result of differences in several interacting parameters. These parameters include ISI, crosstalk, loading, impedance variations, and line length variations.

2.2.2.2 Simulation Parameters

In order to give designs the greatest amount of freedom, a number of assumptions were made. These assumptions were based on our previous simulations, past experience, routing studies, and feedback from some vendors.

For the add-in card, routing studies have shown that a board can be routed with the following requirements. The maximum trace length is 1.5 inches. Data lines must meet a line length skew of ± 0.25 inches with respect to their associated strobes. A trace spacing must be at least 1:3. That is, a 5 mil trace with 15 mil spacing (air gap) or a 6 mil trace with 18 mil spacing is required.

The add-in card impedance is specified as $60\ \Omega \pm 10\%$. In order to achieve the $60\ \Omega$ board requirement, it is recommended that the PCB stack use number 2116 prepreg (nominal 4.5 mil thickness) and the routing be entirely microstrip (outer layers). Signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is desirable to maintain a cleaner and consistent signal return path.

For the motherboard, routing studies have shown that several different requirements exist. For designs that require long line lengths (up to 7.25 inches), a 1:4 spacing is needed. For designs that use shorter line lengths (< 6.0 inches), the amount of trace spacing can be reduced. Simulations should be run to ensure that the design will meet all of the specified requirements.

To get motherboard line lengths of 7.25 inches, a 0.125 inch line length mismatch between the strobe and the data is required. For shorter line lengths, this variation can be relaxed, but simulations should be run to determine by how much.

It is strongly recommended that the motherboard impedance be specified as $60\ \Omega \pm 10\%$. In order to achieve the $60\ \Omega$ board, it is recommended that the PCB stack use number 2116 prepreg and the routing be entirely microstrip. As with the add-in cards, signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is desirable to maintain a cleaner and consistent signal return path.

It is necessary to ensure that all motherboards will be able to run in all transfer modes supported by the system electronics. However, it is expected that an interconnect system designed to support AGP 4X will also support AGP 2X and AGP 1X.

2.2.2.3 Buffers

The models for both the Master and the Target are identical. Both fast and slow buffer models were used in the simulations. Example V-I characteristic curves can be found in Appendix A. In all cases, when a fast buffer was used in our simulations, it was used as both master and slave. The same was also true for the slow buffers.

The buffer models were created from the *AGP Interface Specification, Revision 2.0*. These models also used power supply diodes. The package was modeled with an effective inductance of 10.0 nH.

The models also include an effective capacitance. This capacitance can vary between 3.0 pF and 8.0 pF. For the signal quality simulations, capacitance value of 3.0 pF was used. For the 4X mode simulations, a value of 8.0 pF was used for both the strobe and the data. For flight time simulations, the loading used was 3.0 pF for minimum flight times and 8.0 pF for maximum flight times.

2.2.2.4 Interconnect

Both the motherboard and the add-in card were modeled with a $54\ \Omega$ and $66\ \Omega$ typical transmission line. The $54\ \Omega$ boards had a velocity of 1.71 ns/ft, and the $66\ \Omega$ had a velocity of 1.76 ns/ft. The connector was modeled with a 150 ps delay with two different impedance values, $35\ \Omega$ (slow) and $80\ \Omega$ (fast).

The strobe and strobe# used interconnect impedance values that were based on an odd mode coupling between them. They were also affected by the coupling of the nearest traces. With the spacing between the strobes (strobe and strobe#) being less than the spacing to the nearest traces, the coupling effect is greater between the strobe and strobe#. The nearest data traces could be either odd mode coupled or even mode coupled to the strobes. Odd-odd mode refers to the coupling when the strobe or strobe# is

modeled with the nearest trace in odd mode. Odd-even mode refers to the coupling when the strobe or strobe# is modeled with the nearest trace in even mode.

The data lines were modeled using both odd mode coupling and even mode coupling. Odd mode crosstalk is when the adjacent lines are switching in opposite directions to the trace under analysis. Even mode crosstalk is when the adjacent lines are switching with the trace under analysis.

The stackup used for this analysis used a number 2116 prepreg material that has a 4.5 mil typical thickness. The dimensions used for the typical boards (54 Ω and 66 Ω) are shown below in Table 29. The values used were based on microstrips in a 60 $\Omega \pm 10\%$ printed circuit board (PCB).

Table 32: Board Parameters

Symbol	Parameter	Min	Max	Units	Notes
Er	Dielectric constant	4.3	4.7		1, 5
h	Dielectric thickness	4.1	5.1	mil	2, 5
t	Trace thickness	1.5	2.8	mil	3, 5
w	Trace width	4.8	5.2	mil	4, 5

Notes:

1. Er will probably go as low as 3.9, but at these lower values, the effect will be faster propagation velocity and less variation in the propagation velocity. The effect on impedance is minimal.
2. The maximum thickness of 2116 prepreg is ≈ 5.0 mil.
3. 1 ounce copper has a thickness of ≈ 1.4 mil. The variation comes from plating.
4. Trace width variations are used to control the impedance by PCB vendors. These were varied to get the typical board impedance.
5. Typical 66 Ω board used Er = 4.3, h = 5.1 mil, t = 1.5 mil, and w = 5.2 mil.
Typical 54 Ω board used Er = 4.7, h = 4.1 mil, t = 2.8 mil, and w = 4.8 mil.

Other stackups can be used as long as the combination of the stackup and the trace layout maintains coupled impedance values of 48 Ω to 73 Ω . One such stackup would be a 60 $\Omega \pm 15\%$ board that uses a trace spacing that is 4 times the prepreg thickness. Results from a 2D field solver show that the parameters for this stackup are nearly the same as those for a 60 $\Omega \pm 10\%$ board that uses a trace spacing that is 3 times the prepreg thickness.

Any variations to the recommended stackup and layout need to be simulated. This is to mainly ensure that the coupled impedance values do not exceed 73 Ω , as these may experience ringback failures.

2.2.2.5 Crosstalk

Crosstalk was factored into the simulations using the methodology described in Appendix A. The uncoupled board parameters used were taken from Table 32.

The results for data line type of coupling are shown in Table 33. Data line routing rules in this analysis used 5 mil traces with spacing to the nearest traces of 15 mil and 20 mil.

Table 33: Crosstalk Parameters for 5 mil Data Lines

Trace spacing	Board Z_0	Coupled mode ³	Coupled Z_0	Coupled S_0	Notes
20 mil	54 Ω	odd	50 Ω	1.69 ns/ft	1
		even	57 Ω	1.83 ns/ft	
	66 Ω	odd	63 Ω	1.64 ns/ft	
		even	71 Ω	1.78 ns/ft	
15 mil	54 Ω	odd	48 Ω	1.65 ns/ft	2
		even	58 Ω	1.87 ns/ft	
	66 Ω	odd	60 Ω	1.61 ns/ft	
		even	73 Ω	1.81 ns/ft	

Notes:

1. Required for motherboards to reach maximum line lengths.
2. May be used for shorter line lengths.
3. “Odd mode” and “even mode” refer to the effective impedance realized by a “victim” line at the center of a group of three lines. Odd mode \Rightarrow center line is switching in the opposite direction of the other two lines; even mode \Rightarrow all three lines are switching in the same direction.

The results for strobe line type of coupling are shown in Table 34. The routing for the strobes takes advantage of the strobe and strobe# being in odd mode to each other. The effect of the nearest lines has a lesser amount of coupling effect on the strobes. In the two cases shown, the strobes have either a 15 mil or 20 mil spacing between them. The nearest line in both cases is 30 mil spacing.

Table 34: Crosstalk Parameters for Strobe Lines

Trace spacing	Board Z_0	Coupled mode (strobe/data)	Coupled Z_0	Coupled S_0	Notes
20 mil / 30 mil	54 Ω	odd / odd	51 Ω	1.71 ns/ft	1, 2
		odd / even	52 Ω	1.75 ns/ft	
	66 Ω	odd / odd	64 Ω	1.66 ns/ft	
		odd / even	66 Ω	1.70 ns/ft	
15 mil / 30 mil	54 Ω	odd / odd	50 Ω	1.69 ns/ft	1, 3
		odd / even	51 Ω	1.73 ns/ft	
	66 Ω	odd / odd	62 Ω	1.65 ns/ft	
		odd / even	65 Ω	1.69 ns/ft	

Notes:

1. The first mode in the coupled mode column is the strobe to strobe# coupling. The second mode is the strobe to any other trace coupling.
2. 20 mil spacing from strobe to strobe# and 30 mil spacing to any other trace. Motherboard parameters.
3. 15 mil spacing from strobe to strobe# and 30 mil spacing to any other trace. Add-in card parameters.

2.2.2.6 Flight Time

All flight time simulations were run using 32 different combinations of data signal interconnects. These include four velocity and impedance combinations of motherboard, four velocity and impedance combinations of add-in card, and two impedance combinations of connectors. The four combinations of board interconnect are an uncoupled 54 Ω board in both odd and even mode and an uncoupled 66 Ω board in both odd and even mode. The two combinations of connectors are 35 Ω (odd mode) and 80 Ω (even mode).

Analysis of the minimum flight time showed that this occurs when the add-in card data line length is at the shortest length (0.0 inches) using a “fast” buffer model and 3 pF loading. For these simulations, both the 30 Ω buffer with a curvature of 1.1 and the 38 Ω buffer with a curvature of 2.0 were used. Refer to Section 1.2.2.2 for an explanation of buffer curvature.

Minimum flight time simulations were run with the 32 combinations of interconnect while varying the motherboard line length from 0.0 inches to 2.5 inches in 0.25 inch increments.

Analysis of the maximum flight time showed that this occurs when the add-in card data line length is at the longest length (1.5 inches) using the “slow” buffer model and 8 pF loading. The 48 Ω buffer with a curvature of 2.0 was used for these simulations.

Maximum flight time simulations were run with the 32 combinations of interconnect while varying the motherboard line length from 5.0 inches to 7.5 inches in 0.25 inch increments.




Figure 52: Reference Load

Additional simulations were run to find the reference delay for the buffer models that were used. This value is used to calculate the common clock flight times. (Flight time is defined to be the difference between the delay simulated into the interconnect and the delay simulated into this reference load.) The reference delay is the delay of the driver when driving a reference transmission line load (sometimes referred to as a test load). This is found by driving a 12 inch reference transmission line with an impedance of 60 Ω and a propagation velocity of 1.75 ns/ft. This is pictured in Figure 52. V_{IOUT} is measured by recording the time that the voltage at the open end of the reference transmission line reaches 0.75 V. The reference delay is found by subtracting 1.75 ns from V_{IOUT} . Refer to the AGP Interface Specification, Revision 2.0 Section 4.2.3.3.

All simulations used 1.5 V for V_{oh} and 0.0V for V_{ol} . Also, a threshold value of 0.75 V (0.5 V_{ddq}) was used for all the flight time simulations. Flight time was found by subtracting the reference delay from the time that the signal reached the receiver's threshold.

2.2.2.7 Flight Time Skew

Flight time skew simulations need to simulate any parameter that could cause a skew between two signals. The flight time skew simulations included varying the motherboard and add-in card line lengths, the effective capacitance in the buffer models, and crosstalk on each of the different interconnect combinations. They also included data pattern dependencies, a major contributor to ISI induced skews.

For the source synchronous signals, the various parameters have a cumulative effect on the total flight time skew but do not add linearly. The maximum skew does not occur where each effect is maximized, and the manner in which each effect adds or interacts is not easily predicted. This means that a large number of cross variations need to be looked at.

In all cases, the strobe simulations were compared to the data simulations. Also, all flight time skew simulations were run with the add-in card line length at both 1.25 inches and 1.5 inches. This was done to simulate the worst case strobe-data mismatches contributed by the add-in card. The motherboard line

length was varied from 5.0 inches to 7.5 inches in 0.25 inch increments. This was done to find the maximum line lengths as they relate to skew.

All of the strobe simulations were done using the different types of interconnect combinations for the motherboard (54 Ω and 66 Ω uncoupled board), the connector, and the add-in card (54 Ω and 66 Ω uncoupled board). The motherboard interconnects used coupled parameters for both odd-odd mode and odd-even mode using 1:4 trace:space between the strobos and 1:6 trace:space between the strobos and the nearest trace. The add-in card used coupled parameters for both odd-odd mode and odd-even mode using 1:3 trace:space between the strobos and 1:6 trace:space between the strobos and the nearest trace. The connector used 35 Ω (odd mode) and 80 Ω (even mode).

All of the data simulations were done using the different types of interconnect combinations for the motherboard (54 Ω and 66 Ω uncoupled board), the connector, and the add-in card (54 Ω and 66 Ω uncoupled board). These values are to include the effect of even and odd mode coupling in the simulations. These simulations are grouped by motherboard routing rules. The values used with these rules are shown in Table 33 and Table 34. The add-in card uses the 1:3 spacing, and the motherboard uses 1:4 spacing.

The capacitive loading used for both the data and strobe simulations was 8 pF.

The strobe simulations were run using a 7.5 ns period with a 50% duty cycle time. This is due to the fact that the strobe will constantly be switching states when transferring data. Four cycles were run.

The data simulations were run varying the number of cycles depending on the cycle time. When the duration (halfcycle) was set to 3.75 ns, the number of cycles was set to 4. When the duration was set to 7.5 ns, the number of cycles was set to 2. With the simulator that was used for this analysis, the starting of the first cycle represented a case when the signals were settled. This is due to the symmetry of the buffers and the fact that the initial starting point was set to V_{ss} . The duration was not set to 11.25 ns as shown in Figure 49. Showing an example of ISI with a 11.25 ns duration made the difference clearer because the first rising edge was the same for both simulations.

2.2.2.8 Signal Quality

Signal quality was measured using a different set of simulations, but it is possible to use the results of the flight time and skew simulations to find the worst case results. Overshoot is measured at the highest voltage, and undershoot (negative overshoot) is measured at the lowest voltage at the receiver.

Ringback from the rising edge is the lowest voltage that a receiver comes down to after crossing the threshold at V_{ddq} (1.5 V). Ringback from the falling edge is the highest voltage that a receiver comes up to after crossing the threshold at V_{ss} (0.0 V). Figure 53 is an example that shows how overshoot, undershoot, and ringback are measured.




Figure 53: Signal Quality Measurements

2.2.3 Simulation Results

2.2.3.1 Signal Quality Simulations

The signal quality simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 1.5 inches and a connector line length of 0.9 inches. These simulations used a motherboard line length of 0.0 inches to 7.5 inches in 0.25 inch increments.

The buffer that found the largest amount of ringback was the 38 Ω buffer with a curvature of 2.0 and a slew rate of 2.8 V/ns. The amount used for component loading was 3 pF.

All of the simulations were run with two different cycle times to incorporate the effect of ISI on ringback. These cycle times were 15 ns and 30 ns. The results of these simulations are shown in Figure 54.




Figure 54: Ringback Results

2.2.3.2 Minimum Flight Time Simulations

The minimum flight time simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 0.0 inches and a connector line length of 0.5 inches. These simulations used a motherboard line length of 0.0 inches to 2.5 inches in 0.25 inch increments.

The simulations used two different buffer types. One of the buffers was the 38 Ω buffer with a curvature of 2.0. The other buffer was the 30 Ω buffer with a curvature of 1.1. Both of these buffers were simulated with two different slew rates, 1.4 V/ns and 2.8 V/ns.

The component loading for all of these simulations was 3 pF. Also, the simulations used a 15 ns cycle time.

The results of these simulations show that there are no minimum line length requirements. These results are based on the technique that is described in Section 2.2.2.5. Figure 55 is a graph of the results.




Figure 55: Minimum Flight Time Results

2.2.3.3 Maximum Flight Time Simulations

The maximum flight time simulations were performed on each of the 32 variations of data interconnects. They were also done with an add-in card line length of 1.5 inches and a connector line length of 0.9 inches. These simulations used a motherboard line length of 5.0 inches to 7.5 inches in 0.25 inch increments.

The simulations used two different buffer types. One of the buffers was the 48 Ω buffer with a curvature of 2.0. The other buffer was the 48 Ω buffer with a curvature of 1.1. Both of these buffers were simulated with two different slew rates, 1.4 V/ns and 2.8 V/ns.

The component loading for all of these simulations was 8 pF. Also, the simulations used a 15 ns cycle time.

These simulations assumed a board stackup that consisted of microstrips. Striplines will reduce the maximum line lengths due to reduced velocity of striplines. A 25% reduction in maximum trace lengths when using striplines will ensure that the specified t_{OFF} will be met.

The results of these simulations show that there are no minimum line length requirements. These results are based on the technique described in Section 2.2.2.6. Figure 56 is a graph of the results.




Figure 56: Maximum Flight Time Results

2.2.3.4 4X Transfer Mode Simulations

Analyzing skew requires many more simulations than flight time. This is due to the variations for the strobe interconnect that were compared to each of the data simulations.

A simple calculation of the number of simulations required for each of the interconnect parameters would be 32^2 (1024). These combinations are shown in Table 35. Each of these simulations would be repeated for each of the different buffer types (impedance, curvature, and slew rate) and interconnect line lengths. This is only partly true because there are some short cuts that can be taken to reduce this number.

Table 35: Example of the 1024 Combinations of Interconnect Parameters

Interconnect Parameter	Data		Strobe	
Motherboard Z0	54 Ω / 66 Ω	(2X)	54 Ω / 66 Ω	(2X)
Add-in card Z0	54 Ω / 66 Ω	(2X)	54 Ω / 66 Ω	(2X)
Connector Z0	35 Ω / 80 Ω	(2X)	35 Ω / 80 Ω	(2X)
Motherboard crosstalk mode	Even / odd	(2X)	Odd-even / Odd-odd	(2X)
Add-in card crosstalk mode	Even / odd	(2X)	Odd-even / Odd-odd	(2X)

There are several short cuts that can be taken to reduce the number of simulations. First of all, when the data uses typical impedance for either the motherboard or the add-in card, the same typical impedance is used for the strobe. Another way to reduce the number of simulations is to divide the hold skew simulations from the setup skew simulations. This will allow the hold skew simulations to use only odd mode coupling for the motherboard data and odd-even mode for the strobe. The setup skew simulations will use even mode coupling for the motherboard data and odd-odd mode coupling for the motherboard strobe. This method finds the worst cases with the fewest simulations.

Splitting the simulations up into hold skew and setup skew has added benefits by incorporating line length deltas in the simulations. The hold skew simulations will have the strobe simulated with the maximum amount of line length delta. The setup skew simulations will have the strobe simulated with the minimum amount of line length delta.

Table 37 is an example showing how the data and strobe simulations match up. This is an example of hold skew simulations that use a typical 54 Ω board for both the motherboard and add-in card. This format is repeated using typical 66 Ω boards. Table 37 also shows combinations using a 54 Ω motherboard with a 66 Ω add-in card and 66 Ω motherboard with a 54 Ω add-in card. This puts the number of combinations at 128 when the setup skew and the hold skew are separated. These are shown in Table 36.

Table 36: Reduced Set of Interconnect Parameters to be Simulated

Interconnect Parameter	Data	Strobe
Motherboard Z0	54 Ω / 66 Ω (2X)	Same as data
Add-in card Z0	54 Ω / 66 Ω (2X)	Same as data
Connector Z0	35 Ω / 80 Ω (2X)	35 Ω / 80 Ω (2X)
Motherboard crosstalk mode (setup skew)	Even	Odd-odd
Motherboard crosstalk mode (hold skew)	odd	Odd-even
Add-in card crosstalk mode	Even / odd (2X)	Odd-even / Odd-odd (2X)

The setup skew simulations use different coupling modes for the motherboard. The data lines use odd mode for the motherboard while the strobe uses odd / odd mode coupling.

All of these combinations were simulated at each of the motherboard line lengths and six different buffers. The buffers used are the 30 Ω with a curvature of 1.1, the 38 Ω with a curvature of 2.0, and the 48 Ω with a curvature of 2.0. All three of these buffers used slew rates of 1.4 V/ns and 2.8 V/ns. Appendix C.

Sensitivity Analysis, shows the results of these simulations as 3D charts.

Table 37: Example of Simulation Combinations

DATA					STROBE				
Add-in Card		Connector	Motherboard		Add-in Card		Connector	Motherboard	
Typ bd	Mode		Typ bd	Mode	Typ bd	Mode		Typ bd	Mode
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Even	35 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Odd	35 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Even	80 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / even	35 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / odd	35 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / even	80 Ω	54 Ω	Odd / even
54 Ω	Odd	80 Ω	54 Ω	Even	54 Ω	Odd / odd	80 Ω	54 Ω	Odd / even




Figure 57: Setup Skew Results with 1:4 Trace Spacing

Figure 57 shows the results of setup skew simulations with a motherboard interconnect that used the parameters for 1:4 trace spacing. These simulations used all of the six buffer types. The setup skew simulations used an add-in card line length of 1.25 inches for the data simulations and 1.5 inches for the strobe simulations. The motherboard line lengths used are 5.0 inches to 7.5 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line plus 0.125 inches.




Figure 58: Hold Skew Results with 1:4 Trace Spacing

Figure 58 shows the results of hold skew simulations with a motherboard interconnect that used the parameters for 1:4 trace spacing. These simulations used all of the six buffer types. These hold skew simulations used an add-in card line length of 1.5 inches for the data simulations and 1.25 inches for the strobe simulations. The motherboard line lengths used are 5.0 inches to 7.5 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line minus 0.125 inches.




Figure 59: Setup Skew Results with 1:3 Trace Spacing

Figure 59 shows the results of setup skew simulations with a motherboard interconnect that used the parameters for 1:3 trace spacing. These simulations used all of the six buffer types. The setup skew simulations used an add-in card line length of 1.25 inches for the data simulations and 1.5 inch for the strobe simulations. The motherboard line lengths used are 1.0 inches to 6.0 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line plus 0.5 inches.




Figure 60: Hold Skew Results with 1:3 Trace Spacing

Figure 60 shows the results of hold skew simulations with a motherboard interconnect that used the parameters for 1:3 trace spacing. These simulations used all of the six buffer types. These hold skew simulations used an add-in card line length of 1.5 inches for the data simulations and 1.25 inches for the strobe simulations. The motherboard line lengths used are 1.0 inches to 6.0 inches in 0.25 inch increments for the data lines. The strobe line lengths are the same as the data line minus 0.5 inches.

2.2.4 Design Recommendations

2.2.4.1 Board Materials

The recommended stackup is microstrips on a PCB using a number 2116 prepreg material as shown in Figure 61. Number 2116 prepreg is a common dielectric material available to most PCB vendors. The thickness of this material is nominally 4.5 mils with a maximum thickness of less than 5.0 mils. It is made with a denser fiberglass weave, with less thickness variation in manufacturing than thicker prepreg materials, giving better board impedance tolerances. The thinner dielectric decreases coupling to the signal lines relative to adjacent signals thereby reducing crosstalk effects, which also helps reduce skew due to ISI. Both of these reductions are necessary in making a robust AGP 4X design.

Other prepreg materials may be found that will produce coupled impedance values that are within 48 Ω to 73 Ω . As a general rule, a thicker prepreg material requires a greater amount of spacing between the traces.

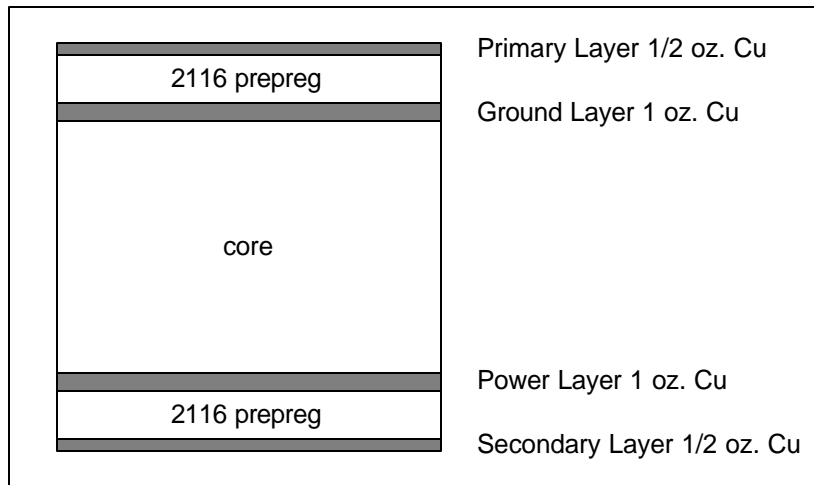


Figure 61: Recommended PCB Stackup

2.2.4.2 Routing Measurement Points

With the tightened requirements for line lengths and line length matching, it is important to define the measurement points. The measurement points for both the upper and lower contact points are on the same horizontal line across the connector. Figure 62 shows where this line is located on the add-in card edge fingers. For the motherboard, the measurement point is at the connector pin.

The measurement points for the master and the target devices are the center of the pad that the ball (BGA package) or pin (QFP package) lands on.




Figure 62: Upper and Lower Contact Measurement Points

2.2.4.3 Escaping Components

Additional design freedom might be given to the board interconnects that use BGA packages. In BGA packages, the package trace lengths are generally longer to the outer row of balls than to the inner row of balls. In these cases, it may be possible to bend the line length delta rule by giving the board interconnects that are routed to the outer row of balls an increase in delta.

When “bending” the line length delta rule, there are a couple of important precautions that need to be taken into consideration. One of these precautions is that the line length rule cannot be violated. Another precaution is that the electrical length of the package needs to be known and not assumed.

This practice has the potential to cause failures if the component package ever changes. Any changes to the package must be accounted for in the interconnect. For this reason, this practice is not recommended, but could be used.

2.2.5 Measurements

This section gives some insight into how to correlate simulations to measurements. The importance of correlation is that ringback and skews are specified at the silicon (pad) where they cannot be physically measured.

2.2.5.1 Ringback

Ringback is not difficult to measure with simulations, but for validation it is. With physical measurements taken at the pin or very near the pin, there needs to be a method to extrapolate these measurements to the specification.

Through simulations, the motherboard designer should find the worst case ringback of the design. This worst case ringback is defined at the silicon pad using worst case package models. Using the same simulations, the point where physical measurements are to be taken needs to be recorded. The physical measurements can be compared to this simulation point.

2.2.5.2 Skew

Skew is the difference between two flight times as defined in the following formula: $t_{\text{SKEW}} = t_{\text{DATA}} - t_{\text{STROBE}}$. t_{DATA} is the flight time of the data, and t_{STROBE} is the flight time of the strobe.

Errors in the skew measurement can be caused by the reflections due to the package when taking flight time measurements at the receiving component pin. This ledge becomes wider when measured farther from the pad. An example of these reflections can be seen in Figure 63. This example shows that the timing that is being measured will vary if this ledge is either raised or lowered. This variation has the potential to artificially increase the skew to over 120 ps.

When making physical measurements with a ledge at or below 0.5 V_{ddq}, it is necessary to make this measurement below the ledge. If one of the measurements is below 0.5 V_{ddq}, the other needs to be made at this same voltage level.




Figure 63: Example of Ledge at the Receiver Pin

2.3 8x Mode Board Design Guidelines

This section describes AGP 8X transfer mode (8X mode). In this section, 8X mode refers to the 8X transfer speed (533MT/s) using AGP3.0 bus electrical signaling (800mV nominal signal swing, 50 Ω parallel terminated, 350mV nominal Vref, etc) as defined in the *AGP3.0 Interface Specification, Revision 1.0*. Board designs that support 8X mode will also meet the 4X and 2X mode layout guidelines.

2.3.1 8X Mode Routing Rules

The guidelines in this section cover motherboard and add-in card printed circuit board (PCB) designs with 8X mode capable controllers. The guidelines are more stringent than 4X and 2X mode guidelines. These guidelines will therefore be applicable to 4X and 2X board designs.

Add-in card designs must follow the “Interconnect Requirements” specified in the *AGP3.0 Interface Specification, Revision 2.0*. These layout rules are reiterated in Table 40.

For motherboard designs, there are three ways that these guidelines can be used:

- Designers who do not have the time or resources to complete a full simulation of their design can route their layout following the rules in Section 2.3.1. These rules preserve enough timing margin to ensure a working design without requiring further simulations. Care must be taken to follow these rules precisely, if no further simulations are to be done.
- Designers wishing to create their own set of layout rules can use the information included in Section 2.3.3. This can be done by comparing the results shown in Section 2.3.3 against Table 39 and applying them to your own design. Simulations on the chosen topology should be done to validate any changes.
- Designers who wish to base their designs on their own simulations can follow the methodology outlined in Section 2.3.2 for additional simulations. This section describes how the simulations and analysis are done to generate these guidelines. This process will yield the most design flexibility but will require extensive simulations.

Section 2.3.4 gives additional design information that can be used with all three methods described above.

Due to the tight interconnect timing budget on 8X mode, more than 1 AGP connector(s) in the signal path will likely violate the timing specification.

2.3.1.1 AGP 8X Mode Design Considerations

Scaling the AGP interconnect from 4X transfer mode to 8X mode requires a two-fold increase in interface speed and a 30% reduction in the interconnect skew allowance. Reductions are required in several of the parameters that make up the total skew budget. An example of this is illustrated Figure 64. Line lengths require closer matching in 8X mode than in 4X mode. The loading variation must be reduced in 8X mode. The crosstalk contribution to interconnect skew must be greatly reduced in order meet the smaller timing budget. This means an increased in the trace-to-trace separation distance on 8X printed circuit boards.

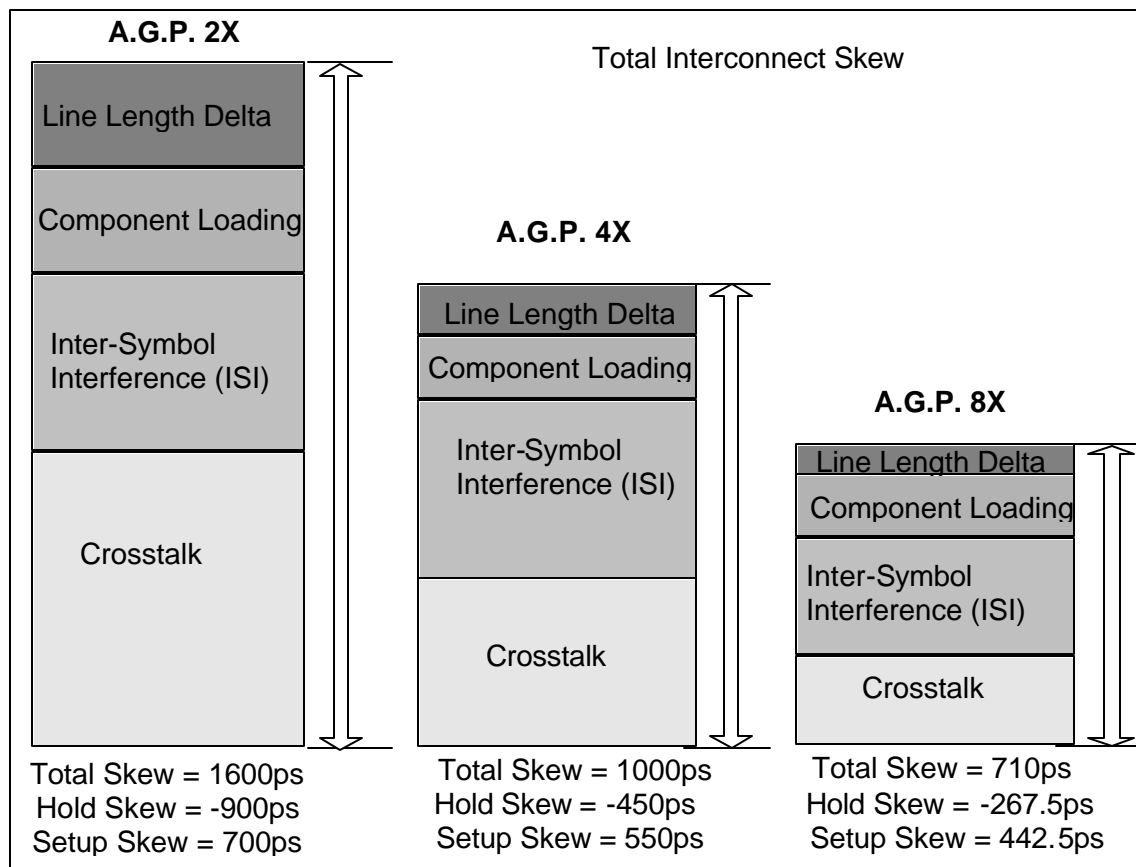


Figure 64: Example of Skew Differences Between 2X, 4X and 8X Mode

Throughout Section 2.3 of this document, the term “data” refers to **AD[31:0]**, **C#/BE[3:0]**, **DBI_HI**, **DBI_LO** and **SAB[7:0]**. The term “strobe” refers to **AD_STBF[1:0]**, **AD_STBS[1:0]**, **SB_STBF** and **SB_STBS**. These are also referred to as the source synchronous signals.

Table 38: 8X Mode Data and Associated 1st and 2nd Strobes

Group	Data	Associated 1 st Strobe	Associated 2 nd Strobe
1	AD[15:0] and C#/BE[1:0]	AD_STBF0	AD_STBS0
2	AD[31:16] and C#/BE[3:2], DBI_HI, DBI_LO	AD_STBF1	AD_STBS1
3	SBA[7:0]	SB_STBF	SB_STBS

2.3.1.2 Board Constraints

Two main factors contribute to PCB impedance variation: manufacturing tolerances and crosstalk. The range over which these factors vary must be reduced to decrease the total impedance variation.

Reducing the board manufacturing tolerances is achieved by specifying the board and layer stackup to board vendors. For AGP 2X mode, the allowable impedance range was 50 Ω to 85 Ω . For 4X mode, the impedance range was reduced to 60 $\Omega \pm 10\%$. For 8X mode, the impedance of the required microstrip traces is specified at 60 $\Omega \pm 10\%$, and the impedance for stripline traces is specified at 56 $\Omega \pm 10\%$.

Reducing the amount of coupling between lines will reduce the impedance variation due to crosstalk. Crosstalk can be reduced by decreasing the distance of a trace to the nearest reference plane (dielectric thickness/height). Crosstalk can also be reduced by increasing the distance to adjacent traces (trace spacing).

In the past it has been common to specify the circuit board layout constraints in terms of the ratio of **trace spacing S** to **trace width W** (e.g. **S/W**). A more useful design constraint depends on the ratio of spacing to **dielectric height H** above a reference plane (e.g. **S/H**). Throughout Section 2.3, trace spacing to height (S/H ratio) is specified. It is the responsibility of the designer to come up with the trace width based on the board stackup to hit the trace impedance target. For example, using a prepreg of dielectric height (H) of 5.0 mils will require a trace spacing (S) of 20 mil for a minimum 4-to-1 S/H ratio.

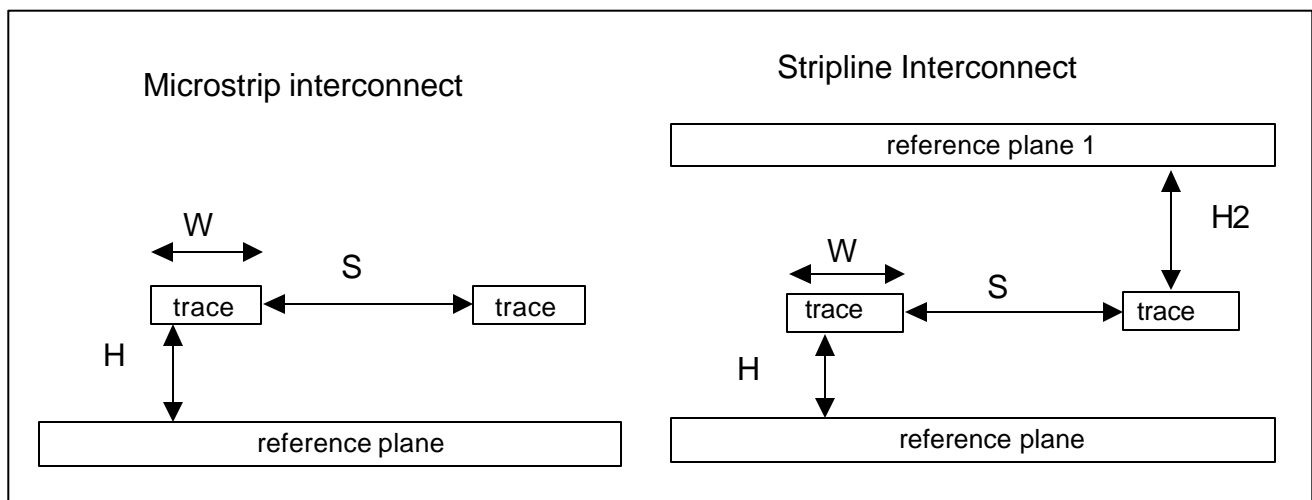


Figure 65: Microstrip and Stripline Trace

The 8X mode source synchronous signals and the 66MHz clock signal should not cross any plane splits (crossing two separate voltage planes). If the common clock signals cross split planes, they should be simulated to ensure signal quality is not compromised.

All 8X mode source synchronous signals should be routed on a layer referenced to a ground plane. To minimize the effect of trace velocity difference between circuit board layers, the source synchronous signals within the same group should be routed on the same layer. The total number of trace via transitions should match within the group. Dummy vias should be used to match the total via count for each group.

2.3.1.3 Source Synchronous Rules for Motherboards

Table 39 gives the routing rules for motherboards. Board layouts following these rules will work without significant additional simulations.

Table 39: 8X Mode Motherboard Rules

Space (S) to Dielectric Height (H) ratio, minimum ¹	Z0	Trace	Physical Line Length (inches)	Line (Electrical) Length Matching ²
5:1	54 Ω to 66 Ω microstrip	1 st Strobe to 2 nd Strobe	2.0 in. \leq line length \leq 6 in.	< 0.005 in. (1.00 ps)
5:1	54 Ω to 66 Ω microstrip	Strobe to Data		
4:1	54 Ω to 66 Ω microstrip	Data	2.0 in. \leq line length \leq 6 in.	Strobe \pm 0.025 in. (5.00 ps) within group
4:1	50 Ω to 62 Ω stripline	1 st Strobe to 2 nd Strobe	2.0 in. \leq line length \leq 7 in.	< 0.005 in. (1.00 ps)
4:1	50 Ω to 62 Ω stripline	Strobe to Data		
4:1	50 Ω to 62 Ω stripline	Data	2.0 in. \leq line length \leq 7 in.	Strobe \pm 0.025 in. (5.00 ps) within group

Notes:

1. These ratios assume a particular board stackup. For different stackups, scale the trace spacing with the dielectric thickness in order to maintain the same spacing to height ratio (S/H). Refer to Section 2.3.4.1 for more information.
2. Line length matching is used as a starting point for the layout. Electrical length matching is required when package trace mismatch is taken into account.

The recommended motherboard impedance range is 60 $\Omega \pm 10\%$ for microstrip traces and 56 $\Omega \pm 10\%$ for stripline traces. This accommodates the commonly used motherboard stackup. The **Space-to-Height (S/H)** ratios in Table 39 specify the minimum trace-to-trace spacing needed to control crosstalk effects. The maximum line length is dependent on the motherboard routing rules. These routing rules allow tradeoffs between signal spacing and line length.

In all cases, it is best to reduce the line length mismatch wherever possible to minimize timing variations. It is also best to separate the traces by as much as possible to reduce trace-to-trace coupling. The physical/electrical length matching requirements include compensation for the package length delta. For motherboard layouts, length matching is specified from pad to connector pins. Refer to Section 2.3.1.7 for more information.

2.3.1.4 Source Synchronous Rules for Add-in Cards

Table 40 gives the routing rules for add-in cards.

Table 40: 8X Mode Add-in Card Rules

Space (S) : Dielectric Height (H) ratio, minimum ¹	Z0	Trace	Physical Line Length (inches)	Line (Electrical) Length Matching ²
5:1	54 Ω to 66 Ω microstrip	1 st Strobe to 2 nd Strobe	1.0 in. \leq line length \leq 1.5 in.	< 0.005 in. (1.00 ps)
5:1	54 Ω to 66 Ω microstrip	Strobe to Data		
4:1	54 Ω to 66 Ω microstrip	Data	1.0 in. \leq line length \leq 1.5 in.	Strobe \pm 0.025 in. (5.00 ps) within group
4:1	50 Ω to 62 Ω stripline	1 st Strobe to 2 nd Strobe	1.0 in. \leq line length \leq 1.5 in.	< 0.005 in. (1.00 ps)
4:1	50 Ω to 62 Ω stripline	Strobe to Data		
4:1	50 Ω to 62 Ω stripline	Data	1.0 in. \leq line length \leq 1.5 in.	Strobe \pm 0.025 in. (5.00 ps) of group

Notes:

- These ratios assume a particular board stackup. For different stackups, scale the trace spacing with the dielectric thickness. Refer to Section 2.3.4.1 for more information,. Dielectric height refers to the distance between the signal trace to the reference plane. Space refers to the trace-to-trace air gap (Figure 65).
- Line length matching is used as a starting point for the layout. Electrical length matching is required when package trace mismatches are taken into account.

The required add-in card impedance range is 60 $\Omega \pm 10\%$ for microstrip and 56 $\Omega \pm 10\%$ stripline traces. This is done to accommodate the commonly used board stackup. The **Space-to-Height (S:H)** ratios in Table 40 specified the minimum trace-to-trace spacing to reduce crosstalk coupling effects.

The physical/electrical length matching requirements include board trace compensation for the package length delta. For the add-in card, length matching is specified from pad to connector edge fingers. Refer to Section 2.3.1.7 for more information.

The maximum line length is the absolute longest physical trace length on the add-in card. With line length matching to compensate package length mismatches, most data traces will be shorter than 1.5 inches on the add-in card. In all cases, it is best to reduce the line length mismatch wherever possible to avoid compromising the timing margin. It is also best to separate the traces by as much as possible to reduce the amount of trace-to-trace coupling.

2.3.1.5 Control Signal and Clock Recommendations

Table 41 gives the routing rules for the control signals and clock lines.

Table 41: 8X Mode Control Signal Line Length Recommendations

Space (S) : Dielectric Height (H) ratio, minimum ¹	Board	Trace	Physical Line Length (inches)
3:1	Motherboard	Control Signals	0.0 < line length < 7.5 in.
4:1	Motherboard	Clock	
3:1	Add-in Card	Control Signals	0.0 < line length < 3.0 in.
4:1	Add-in Card	Clock	4.0 in \pm 0.25 in

Notes:

1. These ratios assume a particular board stackup. For different stackups scale the trace width with the dielectric thickness. Refer to Section 2.3.4.1 for details.

There are no external pull-up or pull-down resistors for control signals on AGP 8X mode board design. The resistors are integrated inside the device buffers.

The clock lines on the motherboard should be simulated to determine their maximum line length. The motherboard must be designed to the type of clock driver that is being used. In all cases, it is recommended to have a trace spacing to any adjacent signals that is at least 4 times the dielectric height.

2.3.1.6 Routing Measurement Points

With tightened requirements for line lengths and line length matching, it is important to define the measurement points. For the add-in card, the measurement point for both the upper and lower contact points are on the same horizontal line across the connector. Figure 66 shows where this line is located on the add-in card edge fingers. For the motherboard, the measurement point is at the connector pin.

The measurement points for the AGP 8X devices are the center of the pins (BGA package) that the balls land on.




Figure 66: Upper and Lower Contact Measurement Points Near Edge Fingers

2.3.1.7 Package Trace Matching

In order to meet the electrical length matching requirements (Table 42, Table 43), the board needs to compensate for the package trace length mismatches.

On the add-in card, the overall electrical length needs to be matched from die pad to the connector edge-fingers. The maximum physical length on the add-in card is 1.5 inches for the source synchronous signals. This corresponds to the shortest package trace. The other source synchronous signals will be less than 1.5 inches physical lengths on the board in order to match the electrical lengths. The package trace versus the circuit board trace velocities needs to be account for when coming up with the overall electrical lengths. Table 42 shows the calculations to achieve matching electrical lengths for traces within the same group.

Table 42: Package Trace Matching Example

Source Sync Signals	Package trace length (in)	Package trace velocity (ps/in)	Board stripline trace velocity (ps/in)	Board trace length (in)	Electrical lengths (ps)
AD0	0.532	167.2	173.5	1.385	329.2
AD1	0.451	167.2	173.5	1.463	329.2
AD_STBF0	0.562	167.2	173.5	1.357	329.4
AD_STBS0	0.413	167.2	173.5	1.500	329.3

On the motherboard, the overall electrical length needs to be matched from die pad to the connector pins. The maximum physical trace on the motherboard is 6 inches (microstrip) and 7 inches (stripline).

2.3.1.8 Ground Plane Reference

Each source synchronous signal group needs to be routed on the same layer referenced to ground with the same layer transitions. Dummy vias are needed to match the total via counts within each source synchronous group near the edge fingers on the add-in card. This is to ensure a clean signal return path to minimize skew between signals within each group.

2.3.2 Simulation Techniques

This section summarizes the interconnect simulations that are done to generate the routing guidelines for the AGP 8X (533 MT/s) interface. The result of these simulations shows that designs using buffers defined in the *AGP3.0 Interface Specification, Revision 1.0* can be done. Figure 67 is a diagram of the topology used in this analysis.

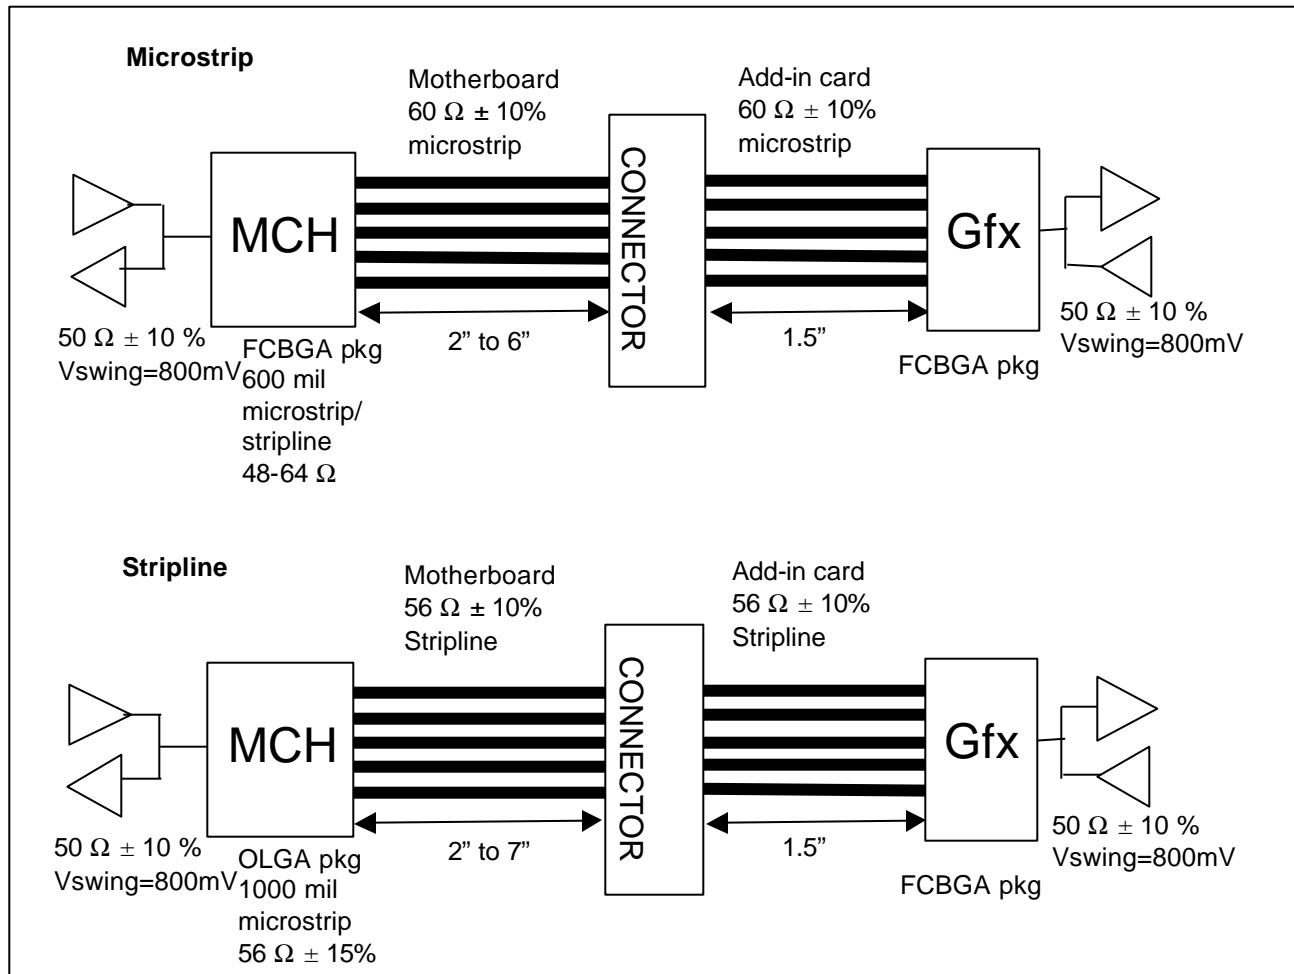


Figure 67: Two Topologies For Interconnect Simulations

The simulation methodology for the AGP 8X is more accurate than the AGP 4X simulations. The simulations for AGP 8X are done with 5-line coupled lossy transmission lines. The AGP4X simulations of Section 2.2 were done with single-line equivalent models using lossless transmission lines. The coupled line simulations can better predict the effect of even and odd mode trace-to-trace coupling to give more accurate timing skew results.

2.3.2.1 Methodology

Three main factors define the range of trace lengths useable for AGP 8X: common clock flight times, source synchronous flight time skews, and signal quality. The common clock flight times do not change from AGP 4X. The maximum line lengths are restricted by flight time skews and signal quality requirements.

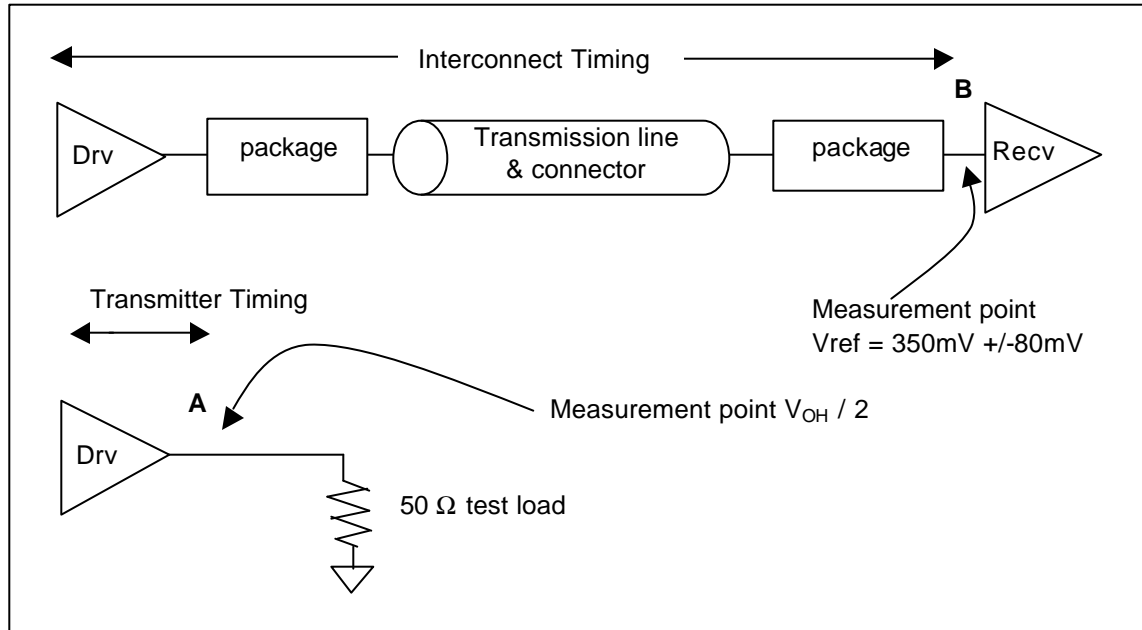


Figure 68: Interconnect Timing Skew Methodology

The interconnect simulation methodology is “**Pad-to-Pad**”. The receiver timing is specified to the pad at point B in Figure 68. The receiver timing measurement point is at $V_{ref} \pm 80\text{mV}$ window. V_{ref} is 350mV and scales with V_{ddq} with a resistor-divider ratio. The transmitter timing is specified at the pad at point A in Figure 68. The transmitter timing measurement point is at $V_{oh} / 2$ terminated to a 50 Ω test load to ground. The interconnect timing for data is the flight time (**T_{data}**) for data to reach B *minus* the transmitter timing for data at A. Similarly, the interconnect timing for strobe is the flight time (**T_{strobe}**) for strobe to reach B *minus* the transmitter timing for strobe at A.

$$\text{Equation 1 : Data interconnect flight time (Tflight_data)} = T_{data@B} - T_{data@A}$$

$$\text{Equation 2: Strobe interconnect flight time (Tflight_strobe)} = T_{strobe@B} - T_{strobe@A}$$

Thus, the interconnect skew timing is the Data *minus* Strobe interconnect flight time.

$$\text{Equation 3: Interconnect skew timing} = T_{flight_data} - T_{flight_strobe}$$

This interconnect skew methodology will include V_{oh} variations at the transmitter. It also includes the $V_{ref} \pm 80\text{mV}$ timing window at the receiver. V_{ref} varies from 320mV to 380mV due to $V_{ddq} \pm 5\%$ and resistor divider accuracy (Figure 69). An additional $\pm 50\text{mV}$ is added on top of that to account for receiver bias mismatch for data/strobe (30mV) and noise on V_{ref} signal due to coupling (20mV). Therefore, a total window of $V_{ref} \pm 80\text{mV}$ (270mV to 430mV) is used for interconnect simulation analysis. Since the simulation is done with a 5-line coupled package model, it will also include the package coupling effects.

This methodology assumes all Tco mismatches, slew rate differences between different data and strobe signals, and any timing skew up to the test load (measured at $V_{oh} / 2$) are included in the transmitter spec timing.

The nominal value of V_{ref} is chosen to be at 350mV and scales with V_{ddq} . The V_{ref} is not at exactly half (400mV) of the target V_{swing} (800mV). This is to balance the overall timing budget in order to offset the naturally biased receiver setup/hold timing. The receiver requires less setup time and more hold time (Table 43). By moving the V_{ref} down to 350mV, the geometric effects will shift the interconnect skew from hold to setup. This will result in a larger interconnect setup skew and a smaller interconnect hold skew. Thus the asymmetric V_{ref} will help to meet the tight timing budget (Table 43). This will only work because data is triggered on rising-edge strobe only (Figure 69). The receiver timing window is chosen at $V_{ref} \pm 80mV$ to cover V_{ddq} variations and noise on V_{ref} signal.




Figure 69: Asymmetric V_{ref} and Receiver Timing Windows

Table 43: 8X Mode Timing Budget

Element	Skew Component	Setup	Hold	Units
Half Bit Time		937.5	937.5	ps
Transmitter	Total skew (refer to Table 6 for details)	410	460	ps
Receiver	Total skew (refer to Table 6 for details)	85	210	ps
Interconnect	ISI, crosstalk (package coupling, trace impedance mismatch)	432	257	ps
	Transmission line length mismatch	10.5	10.5	ps
	Total skew	442.5	267.5	ps
Remaining Margin		0.0	0.0	ps

2.3.2.2 Simulation Parameters

Several assumptions were made to give designers the greatest amount of freedom. These assumptions are based on previous simulations, past experience, routing studies, and feedback from AGP interface design vendors.

For the add-in card, routing studies have shown that a board can be routed with a maximum physical trace length of 1.5 inches for the source synchronous groups. Data line layout must compensate for package length mismatches as described in section 2.3.1.7. Data lines must meet a line length skew of ± 0.025 inches (5.00ps) with respect to their associated strobes. Trace spacing to dielectric height ratio of at least 4-to-1 is required.

In order to achieve the required board impedance, the circuit board stackup should use number 2116 prepreg (nominal 4.5 mil thickness) for designs that require microstrip (outer layers) routing. Signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is required to maintain a clean and consistent signal return path.

Similarly, the motherboard impedance is specified as $60\ \Omega \pm 10\%$ (microstrip) and $56\ \Omega \pm 10\%$ (stripline). As with the add-in cards, signals within a strobe/data group should be routed on the same layer and should not cross breaks in the power/ground plane over which they are routed. Routing all signals over the ground plane is required to maintain a clean and consistent signal return path.

It is necessary to ensure that all motherboards will be able to run in all transfer modes supported by the system electronics. However, it is expected that an interconnect system designed to support AGP 8X will also support AGP 4X due to the tighter layout rules.

2.3.2.3 Buffers

IBIS models are used in simulations with HSPICE. The models for both the MCH and the Gfx devices are identical. The models are generated from the V-I spec curves in the *AGP3.0 Specification, Revision 1.0*. The target buffer impedance (Z_{targ}) is $50\ \Omega$. The device pulldown impedance (Z_{pd}) is $50\ \Omega \pm 10\%$. The min corner corresponds to Z_{pd} of $45\ \Omega$ and the max corner for $55\ \Omega$. The device pullup impedance (Z_{pu}) is specified at $43.75\ \Omega \pm 10\%$ with nominal V_{ddq} (1.5V) and V_{swing} target (800mV). With $V_{ddq} \pm 5\%$ and resistor divider accuracy, Z_{pu} ranges from $38.19\ \Omega$ to $49.78\ \Omega$. This allows the V_{oh} to vary from 714 – 893mV. In addition to 2 corners for different buffer impedance, there are 2 separate buffers with 2.0V/ns and 3.5V/ns slew rate. The models also include an effective pad capacitance of 2.5pF. Example V-I characteristic curves can be found in Appendix B.

2.3.2.4 Packages

The packages for MCH used in simulations are FCBGA (flip-chip or C4) and OLGA. The models are 5-line coupled models. The coupling effects of even and odd mode aggressors toggling are considered in the HSPICE simulations.

The packages for Gfx used in simulations are OLGA and FCBGA flip chip. The flip-chip packages are also 5-line coupled models.

2.3.2.5 Connectors

The AGP connector model is a 4x4 coupled matrix provided by AMP Corporation. The victim and aggressors are chosen corresponding to the actual pinout. Table 44 shows the mapping of the connector model to the 5-line PCB simulation model. For data, "DATA C" is the victim, and the other DATA pins are aggressors. For strobe, "STBS" and "STBF" are the victims, and the other DATA pins are aggressors.

Table 44: AGP Connector Simulation Model Pin Map

B		A		Pin
GND		GND		49
	DUMMY4		DUMMY1	50
DUMMY5		DATA E		51
	VDDQ		VDDQ	52
DATA B		DATA C		53
	DATA A		DATA D	54
GND		GND		55
	DUMMY3		DUMMY2	56

B		A		Pin
GND		GND		55
	DUMMY1		DUMMY3	56
DUMMY2		DATA E		57
	VDDQ		VDDQ	58
STBF		STBS		59
	DATA D		DATA A	60
GND		GND		61
	DUMMY4		DUMMY5	62

2.3.2.6 Printed Circuit Board

Both the motherboard and the add-in card PCB are modeled as W-element lossy transmission lines in HSPICE. The PCB models are 5-line coupled. Both the microstrip and stripline are modeled. The W-element RLGC file is generated from a worst-case field solver sweep based on the corresponding microstrip or the stripline stackup.

2.3.2.7 Crosstalk

Crosstalk is being considered with the 5-line simulation models. The center line is the victim, and the adjacent 2 lines on each side of the victim are the aggressors. The victim data pattern is selected to simulate Inter-Symbol Interference (ISI) effects on timing skews. To simulate the worst case ISI, the aggressors are either switching at the same edge as the victim (even mode coupling) or at the opposite edge as the victim (odd mode coupling). To simulate the worst case coupling to victim, the aggressors are switching at the fastest toggle rate. Some examples of victim and aggressor data patterns are shown in Figure 70.

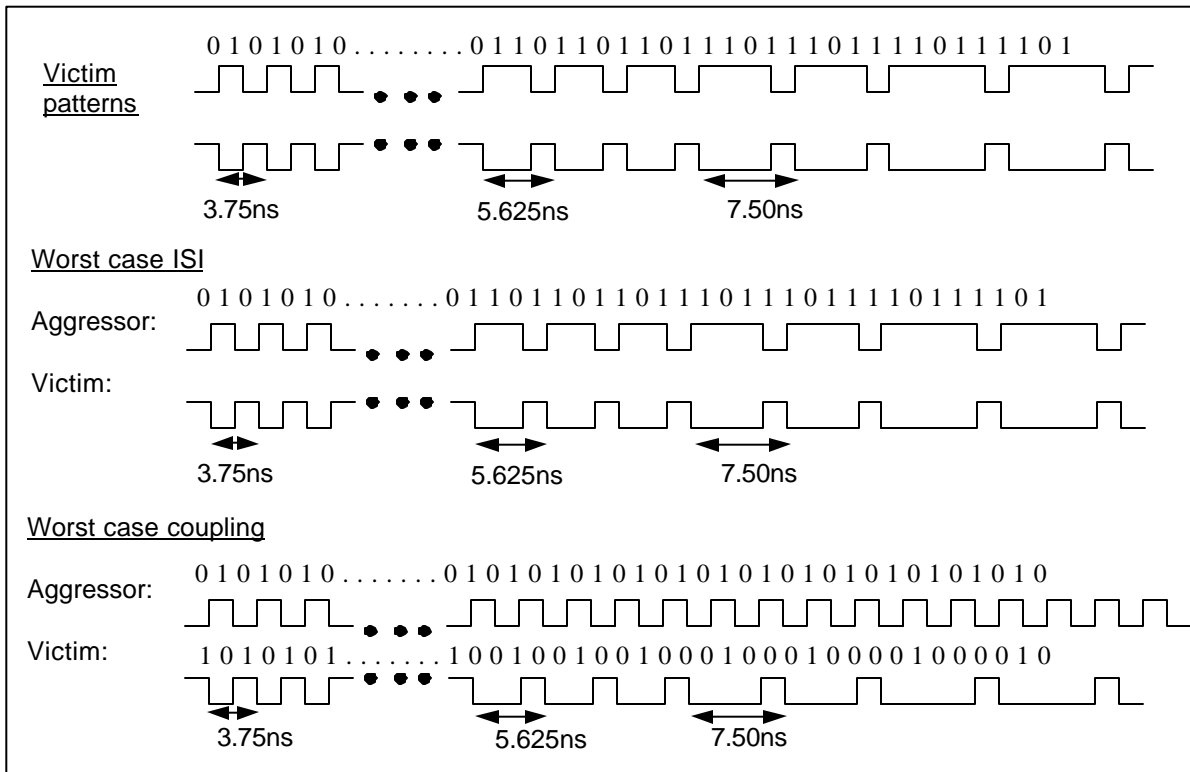


Figure 70: Data Pattern For Simulations

Data line routing rules in all simulations analysis used nominal 5 mil traces with spacing to the nearest traces of 20 mil on a 5 mil thick dielectric. For strobes, the routing rules are 5 mil trace width with 25 mil edge-to-edge spacing for microstrip case.




Figure 71: Data and Strobe crosstalk 5-line model example

The AGP3.0 Specification, Revision 1.0 allows for a short breakout region (500 mil for motherboard chipset and 300 mil for add-in card chipset) with trace spacing tighter than the main route. Figure 72 shows an example of a typical breakout from a FCBGA package.




Figure 72: Trace Breakout Example

The 5-line simulation model is modified for the breakout region as shown in Figure 73.




Figure 73: Data and Strobe crosstalk 5-line model example for breakout region

2.3.2.8 Sweep Matrix

The following is a table for the simulation sweep matrix for all the parameters considered. There are a total of 20 variables in these models. Therefore, a 20 x 20 matrix is created for the entire simulation sweep. Section 2.3.2.9 will describe the criteria for comparing the data to strobe skew based on the associated corners.

Table 45: Simulation Sweep Matrix

Models	Variations			
Driver	2V/ns	2V/ns	3.5V/ns	3.5V/ns
	Min impedance	Max impedance	Min impedance	Max impedance
Receiver	Min impedance		Max impedance	
Package	Min impedance		Max impedance	
Motherboard PCB	Stripline	Stripline	Microstrip	Microstrip
	Min impedance	Max impedance	Min impedance	Max impedance
Add-in Card PCB	Stripline	Stripline	Microstrip	Microstrip
	Min impedance	Max impedance	Min impedance	Max impedance
Vddq	High (+ 5%)		Low (- 5%)	
Crosstalk	Odd		Even	

2.3.2.9 Skew Calculations

Skew calculations are based on simulation data and post processing through scripts and spreadsheet calculation. There are two steps to calculate interconnect skew. First, the transmitter timing must be obtained and subtracted from the flight time. The measurement point for transmitter timing is $V_{oh} / 2$. Second, the interconnect skew is determined by subtracting data and strobe flight values after the transmitter timing calculation. These two steps and their equations are described in Section 2.3.2.1.

Since the strobes on AGP 8X are rising-edge triggered only, the calculation spreadsheet must be able to associate the appropriate strobe (1st or 2nd) and strobe edge with each data cycle. Other skew calculation criteria is given in Table 46.

Table 46: Skew Calculation Criteria

Associated Corners	Explanations
Assume Same Vddq for Driver and Receiver	This means that we should compare simulations with driver and receiver generated at the same Vddq.
Consider receiver switching variations	A $\pm 80\text{mV}$ receiver switching variation should be included in skew calculation. (Figure 69)
Consider reference voltage (Vref) generation variation due to resister divider error	Since Reference voltages (Vref) are generated from resister dividers, one must take into account the uncertainties (included in $\pm 80\text{mV}$) caused by the inaccuracies in these dividers (Figure 69)
Assume same physical parameters when comparing data and strobe	Strobe and data buffers should come from the same corner of simulation model when calculating skew. This means we should compare strobe and data timing at the same silicon and board corner only.
Assume same receiver reference points for data and strobe	Strobe and data measurement voltage should be the same for the skew calculation. Given that receivers will have switching variations, the worst-case skew will always occur when data and strobe switch at the same extremes of the switching region. In AGP8X the receiver switching region extremes are at 430mV and 270mV. (Figure 69)

Consider crosstalk cases	Opposite (even and odd) crosstalk switching modes must be considered in strobe to data skew calculation.
Consider different data edge to strobe edge comparison	Include as many different switching edge comparisons as possible. One should consider shifting simulated data edges with respect to the strobes. See Figure 74 for illustration. In the non-shifting case, data edge “D1” is compared with strobe edge “S1” and so on. Then all the data edges are shifted by 1 cycle ahead. Now the data edge “D2” is compared with strobe edge “S1” and “D3” with “S2” instead. This will generate the worst-case skew by comparing all data edges with all strobe edges across the length of the simulation pattern.

There may be other platform specific parameters that may cause additional skews. The designers must consider them as a part of the skew calculation.

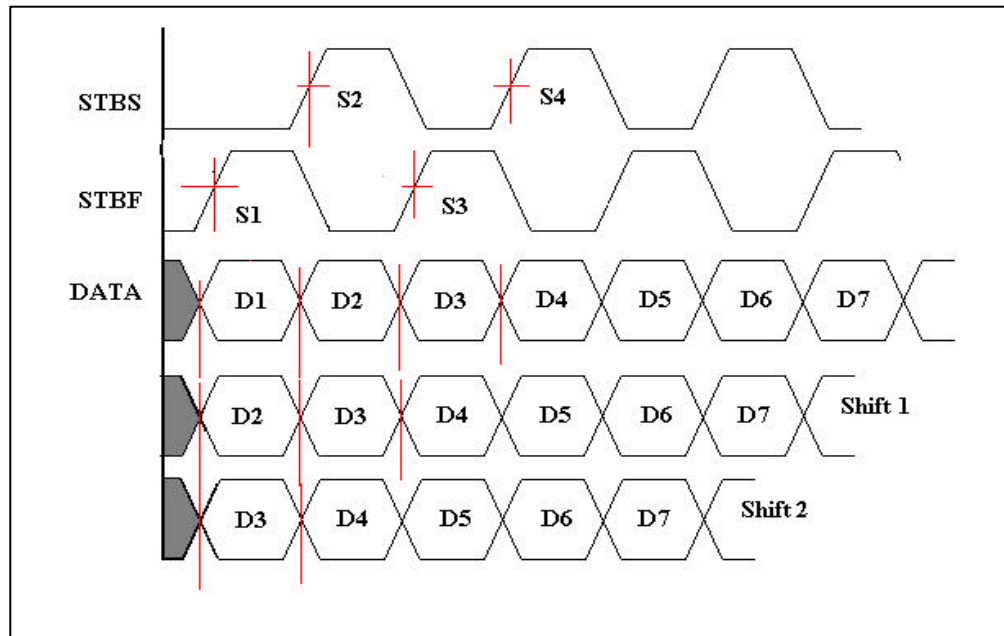


Figure 74: Shifting Data Edges

2.3.2.10 Signal Quality

Signal quality at the receiver is measured using the same set of simulations. Overshoot is measured at the highest voltage (V_{OSH_P} positive), and undershoot (V_{OSH_N} negative) is measured at the lowest voltage at the receiver. Ringback from the rising edge is the lowest voltage that a receiver comes down to the receiver ringback window ($V_{REF} \pm V_{RB}$). Ringback from the falling edge is the highest voltage that a receiver input signal comes up to after crossing the threshold. Ringback is not allowed in the ringback window (the keep out region) until after t_{DRIVE} . Figure 75 shows how overshoot, undershoot, and ringback are measured.

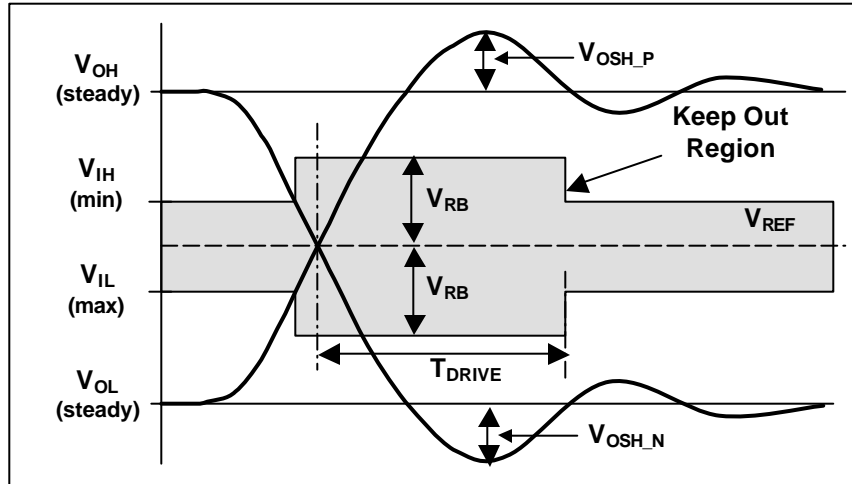


Figure 75: Signal Quality Spec

The waveforms in Figure 76 are used to validate that the input buffers can meet V_{RB} and t_{DRIVE} . Figure 76a shows the rising edge case where the buffer input level is switched from $V_{OL}(\text{nominal})$ to $V_{REF} + V_{RB}$ at the minimum and maximum allowed input slew rates, and over the appropriate temperature, voltage and process corners. The input level is maintained at that level for t_{DRIVE} time and then reduced to $V_{IH}(\text{max})$ at the same slew rate. The speed of the input buffer must not change for this waveform versus an input waveform of greater sustained input high level. The falling edge must also be checked as shown in Figure 76b.




Figure 76: t_{DRIVE} test waveforms

The input buffer should also be characterized using variations on these waveforms to see how sensitive the buffer is to input conditions. The first variation is to vary the high drive input level from $V_{REF} + 0.3$ to V_{REF} volts, without the step to $V_{IH}(\text{Max})$ as shown in Figure 77, and plot the input buffer delay at each step. A characteristic like the one shown in Figure 78 should be seen. The buffer delay should become essentially constant at a voltage less than $V_{REF} + V_{RB}$. Both rising and falling edges and minimum and maximum slew rates must be checked.




Figure 77: Input drive level test waveform




Figure 78: Input delay versus input drive level above Vref

The second variation on the input waveforms, shown in Figure 79, is to step the input drive of the rising waveform to $V_{REF} + V_{RB}$ and vary the time before the input level returns to $V_{IH}(\text{max})$, the “ V_{RB} drive time.” The input buffer delay should vary like the waveform of Figure 80.




Figure 79: Input drive level test waveform




Figure 80: Input delay versus V_{RB} drive time

In this case, the delay should become essentially constant at a time point less than the t_{DRIVE} spec. Again, both rising and falling edges at minimum and maximum slew rates should be checked.

2.3.3 Simulation Results

2.3.3.1 Signal Quality Simulations

The specified limit for overshoot is 350mV above the Voh (steady). The spec limit for undershoot is 300mV below Vss.




Figure 81: Worst Case Overshoot Results for Microstrip and Stripline Topologies




Figure 82: Worst Case Undershoot Results for Microstrip and Stripline Topologies

The specified limit for ringback free region is $V_{ref} \pm 150\text{mV}$. Please refer to section 2.3.2.10 for details on the ringback analysis.




Figure 83: Ringback Results for Microstrip Topology




Figure 84: Ringback Results for Stripline Topology

2.3.3.2 Timing Skew Simulations

The specified limit for worst-case setup skew is 432ps (Table 43). The spec limit for worst-case hold skew is 257ps (Table 43). The worst case timing skews with add-in card length 1.5" and different motherboard line lengths are shown below:




Figure 85: Timing Skew (Setup) For Microstrip and Stripline Topologies




Figure 86: Timing Skew (Hold) For Microstrip and Stripline Topologies

2.3.4 Additional Design Recommendations

2.3.4.1 Board Materials

The recommended stackup is using microstrip on a PCB using a number 2116 prepreg material as shown in Figure 87. Number 2116 prepreg is a common dielectric material available to most PCB vendors. The thickness of this material is nominally 4.5 mils with a maximum thickness of less than 5.0 mils. It is made with a denser fiberglass weave with less thickness variation in manufacturing than thicker prepreg materials, giving better board impedance tolerances. The thinner dielectric decreases coupling to the signal lines relative to adjacent signals thereby reducing crosstalk effects which also helps reduce skew due to ISI. Both of these reductions are necessary in making a robust AGP 8X design.

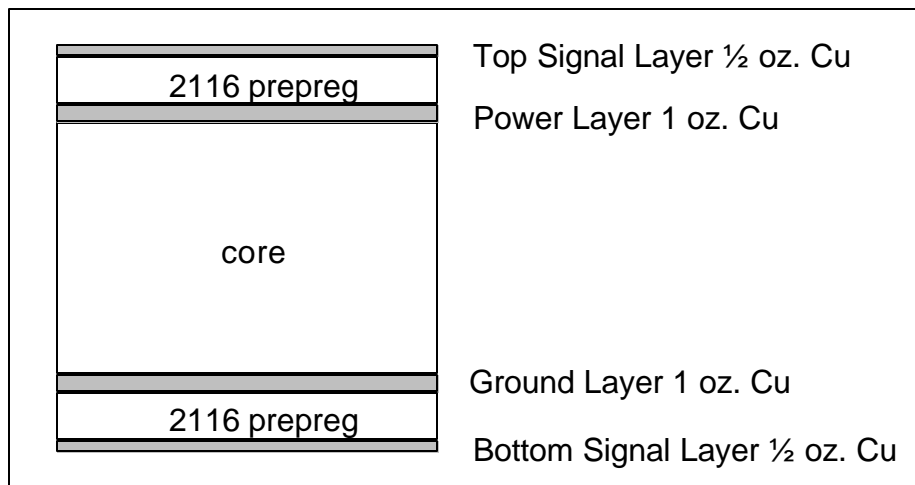


Figure 87: Recommended 4-Layer PCB Stackup

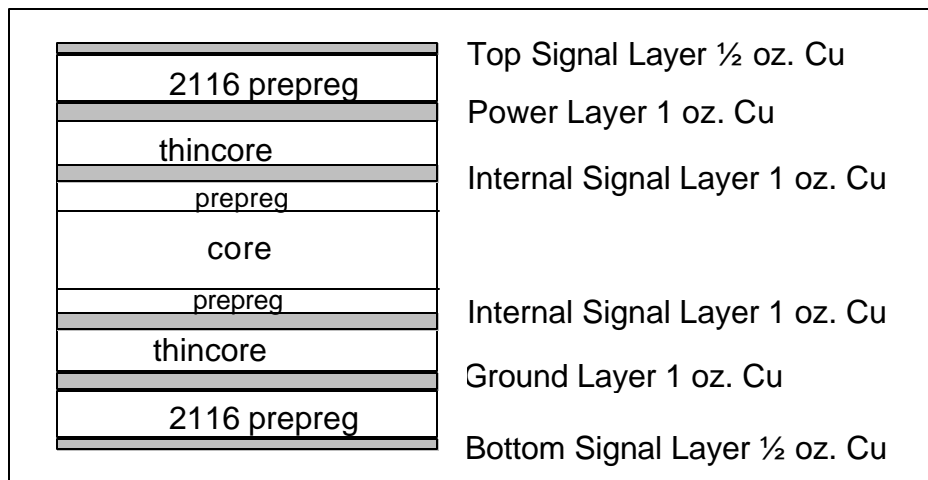


Figure 88: Recommended 6-Layer PCB Stackup

2.3.4.2 VDDQ Plane

The Vddq plane should be generated from a single linear voltage regulator on the motherboard. The plane from the voltage regulator source to the AGP connector Vddq pins should be wide enough to minimize the inductance path. The AGP 8X device buffer I/O rail on the add-in card should connect to the Vddq plane directly. This is to ensure the Vddq rail is shared between the AGP 8X device on the motherboard and the device on the add-in card.

2.3.4.3 Connector Impedance Matching

Typical AGP connector equivalent inductance is about 7 to 8 nH. The equivalent capacitance for the pin and solder tab is about 1 to 2 pF. Therefore the equivalent impedance of the connector is roughly 59 to 89 Ω . The motherboard designer can put short trace stubs near each connector pin to increase the equivalent capacitance by up to 0.5 pF. This will reduce the AGP connector impedance, bringing it closer to the trace impedance of 56 to 60 Ω . This will also help to minimize impedance discontinuities at the connector.

2.3.4.4 Package Pinout

There should be ground pins within the AGP 8X signal field in a ball grid array (BGA) pinout. Each signal should be adjacent to a ground pin for a 2-to-1 signal to ground ratio. It is recommended that the AGP 8X source synchronous signal pins to be assigned to the inner rows. This allows signals to escape on the circuit board from vias to the inner or bottom layers. The common clock signals can be assigned to the first, second, or third row of the BGA since there are no via matching requirements on those signals. The designers should optimize the package pinout with PCB layout taken into consideration.

2.3.4.5 Trace Bends and Serpentes

Trace bends and serpentes will induce a certain amount of trace self-coupling causing the line to be shorter electrically than physically. As a general rule, try to avoid 90 degree bends and use two 45 degree bends instead. The total number of turns on a trace should be limited to less than 15.

2.3.4.6 Measurements

This section explains one possible way how simulation-to-measurement correlation can be performed. The importance of correlation is that Signal Integrity (SI) and Timing parameters are specified at the silicon (pad) while physical measurements normally are done on pin and/or via which prevents direct comparison between simulated and measured numbers.

Explanation in this section will focus on system-level correlation, which implies that findings cannot be generalized and correlation has to be performed for each and all motherboard/AGP card (chipset silicon/graphic silicon) combinations of interest.

General silicon characterization (independent of chipset silicon/graphic silicon combination) usually is performed on tester. Additional information on this characterization can be found in section 3.1.3 of this document.

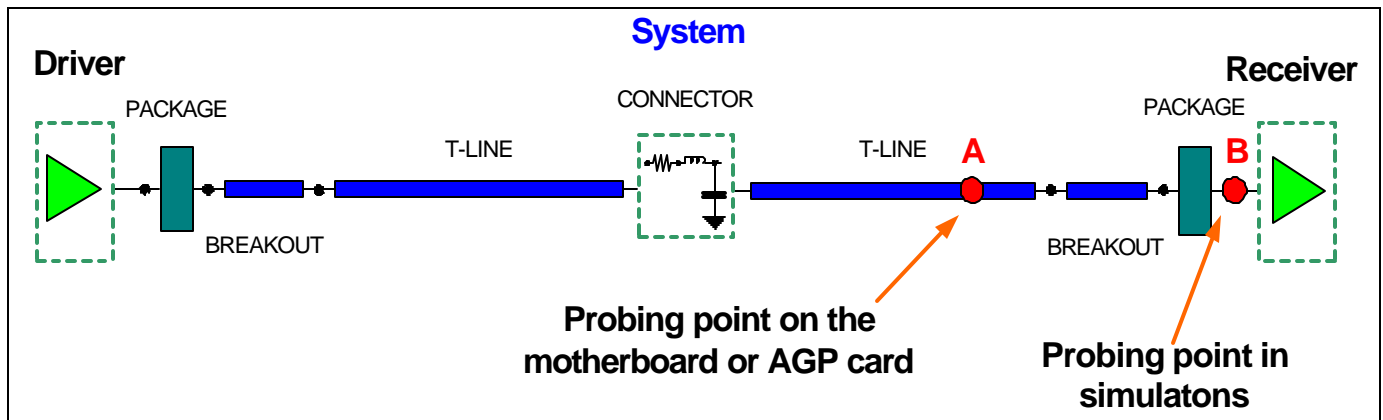


Figure 89: Measurement point definition

2.3.4.6.1 Signal Integrity (SI)

Through post-layout simulations, the motherboard designer should identify the worst Signal Integrity case for a particular design, measured at the pad (point B on Figure 89). During the same simulation runs, waveforms should be captured at the pad (point B) and at the point where physical measurements are to be taken (point A on Figure 89). The point simulated waveform at point B can be compared to the measurements and a correlation can be made to infer the pad waveforms. Usually, the worst case signal conditions are generated by using worst case package models with biggest mismatch between the motherboard and package impedance.

2.3.4.6.2 Skew

Skew is defined as the difference between two flight times as defined in the following formula: $t_{\text{SKEW}} = t_{\text{DATA}} - t_{\text{STROBE}}$. t_{DATA} is the flight time of the data from source pad to the measurement point at the receiver, and t_{STROBE} is the flight time of the strobe from source pad to the measurement point at the receiver.

Whenever it is possible, the approach used for measuring signal integrity should be also employed for measuring skew timings (skew). Non-monotonic edges (ledges, etc.) can generate skew measurement errors. Ledges are created by the reflections due to the package when taking flight time measurements at the receiving component pin. This ledge becomes wider when measurements are performed farther from the pad. An example of these reflections can be seen in Figure 90. This example shows that the timing that is being measured will vary if this ledge is either raised or lowered. This variation has the potential to artificially change the skew.

When making physical measurements with a ledge at or below V_{ref} , it is necessary to make this measurement below the ledge. Similarly, for high-to-low transition, with a ledge at or above V_{ref} , it is necessary to make this measurement above the ledge.

Both strobe and the data measurements need to be done at the same threshold level.




Figure 90: Example of ledge at the receiver pin

Additional correction for skews measured at component pin can be done by including package differences between the data and associated strobe(s). If receiver component package electrical length for data and strobe is respectively $t_{\text{DATA-P}}$ and $t_{\text{STROBE-P}}$, corrected flight time would be $t_{\text{SKEW}} = (t_{\text{DATA}} + t_{\text{DATA-P}}) - (t_{\text{STROBE}} + t_{\text{STROBE-P}})$. If ledges are present, it is necessary to adjust threshold voltages for the flight time measurements, as it is explained above.

Electrical length of the receiver component package represents time-delay associated to the receiver package and it can be measured by the Time Domain Reflectometry (TDR). More details on package TDR measurements can be found in section 3.1.3 of this document.

3 Manufacturing Test and System Validation Guidelines

With increased data rates and demand on interface performance the challenges are not exclusively design related. The ability to screen out defective product and to validate silicon in the system environment becomes more difficult. Understanding the challenges while designing the silicon permits the product team to: add design for testability (DFT) capabilities, plan a test strategy and prepare for system validation.

3.1 Manufacturing Test Overview

A number of manufacturing test challenges exist with AGP8X AC timings that didn't impact the testing of AGP4X AC timings. These challenges include:

- Tight AC timings specs: $T_{va} = 527$ ps; $T_{vb} = 477.5$ ps; $T_{set-up} = 85$ ps; $T_{hold} = 210$ ps;
- Source Synchronous timings;
- Specifying of timings at the pad;
- 533 MT/s data rate.

The purpose of the information in this section is to make design and test engineers aware of the challenges and possible solutions that will enable them to successfully ramp their product to high volume manufacturing (HVM).

3.1.1 Test Platform limitations with Edge Placement Accuracy

Edge Placement Accuracy (EPA) describes the accuracy to which a tester output strobe or input timing edge can be placed within a tester clock period. This accuracy can be broken down into a purely random component and a systematic component. An EPA of 100 ps translates to the fact that the tester strobe/input edge will be placed within ± 100 ps of the desired location. Automatic test equipment (ATE) vendors calibrate a delivered system within the specified total EPA. This section explains the effect EPA can have on tester settings for specified timings and suggests ways to mitigate their effects in order to meet a product's yield loss and escape rate goals.

illustrates EPA for two pins on the device under test (DUT), Clk and Data pin Z34. Since the timing relationship between the two pins is affected by the tester strobe placement on both pins, this fact needs to be accounted for in the setting of tester-programmed limits. The result is illustrated in Figure 92 which shows an example distribution of actual component performance relative to spec. Most parts fall well within the spec with a small tail of the distribution outside spec. These are the units which testing needs to catch. If we placed the test strobe at the specified timings, the test strobe timing uncertainty will cause some parts to fail even though they meet the timing specifications (yield risk area in the graph). Some parts will pass the test even though they fail to meet the timing specifications. When a unit that should have been failed passes the test, it is referred to as an "escape." The rate at which bad parts escape detection is usually measured as defects per million (DPM) and is indicated as the arrow to the right of the spec limit in Figure 92.

A common means of avoiding escapes is to tighten the timing of the tester strobe by 2 EPA as shown in Figure 93. This puts all the test strobe timing uncertainty within the test limit so that there are no test escapes. However, this will affect the yield of the product and, depending upon the silicon timing distribution, this may result in excessive yield loss. The designer could design to a timing target that is 2 EPA tighter than the specified timing. However, this translates into over design of the silicon buffer and may not be feasible in the available silicon process. The design and tester guardbands in Tables 2-5 are not generally large enough to cover 2 EPA.

One solution is to play with the tester strobe position to balance the traditional trade-offs of yield losses and escapes. The solution turns into an empirical analysis over a sample of manufacturing product; but with the increased performance, this may not be a cost effective solution.




Figure 91: Definition of Edge Placement Accuracy




Figure 92: EPA Effect On Testing to Specifications




Figure 93: Moving Tester Strobe to Reduce Escapes

An alternative approach can be derived if we consider the two components of EPA: systematic and random. We can calibrate the test equipment such that all the systematic inaccuracy is understood and accounted for. Then we are left with only the random component. This component is usually much smaller and we can reverse guardband for this inaccuracy (Figure 94). We will detect parts that violate timing specifications if we test the interface for multiple cycles (e.g. 500 cycles). The test and design guardband allowances in Tables 2-5 are more in line with this testing methodology. Calibrating your test equipment further than the initial EPA will take time, but has been shown to be very effective for testing the specifications of high-speed interfaces. Also ATE vendors are trying these kinds of improvements as they seek to improve the accuracy of their test equipment.¹

¹ See Sartori L. and West B.G., "The Path to One Picosecond Accuracy," 2000 ITC Proceedings, pp 619-627




Figure 94: Settings for Random Measurement Error

3.1.2 Accurately Assessing Source Synchronous Timings

Source synchronous (SS) timings deal with the difference in delay (skew) between two signals. The absolute delay from a common timing point does not affect their timing relationship. However, using traditional AC testing methods, changes in the timing components common to both signals affects the measurement. This results in a poor assessment of the timings. Two areas will be addressed in this section: accurate assessment of the differential timings during silicon design validation and HVM testing of these measurements.




Figure 95: Uncertainty When Using Worst-case Timing

The results from a common measuring technique that is used to “measure” the skew between the Strobe and Data are shown in Figure 95. This method uses the minimum difference of the Strobe Tco and the Data Tco (for the spec TDVb, it is $\text{Strobe_Tco_min} - \text{Data_Tco_max}$). Since this method uses the worst-case extremes over all of the vectors within the pattern suite, there is a potential for it to be pessimistic.

As illustrated in Figure 95, the difference between the “Tco Delta” and “Best Case” Tvb is the uncertainty in the result. This uncertainty results from the variation in the position of the output of the strobe signal (jitter) over the vectors within the pattern suite. At a given Vcc and temperature, a jitter on the order of 400 ps is typical.




Figure 96: Example of Source Synchronous Data Timings

The location of the Data and Strobe signals will vary from cycle to cycle. Each signal will have “common jitter” that doesn’t affect the relative delay between signals and “differential jitter” which directly affects the timing. This means that in order to measure Tdvb and Tdva parameters of every one of the SS signals, it is required to measure the Data signal’s location relative to the associated Strobe signal’s location in the same cycle, and then compile the results for every cycle. Then, the worst-case Tdva & Tdvb can be found from all the data that was collected for that particular test vector pattern. The challenge in Design Validation is to overcome the pessimistic results obtained using traditional searches for worst-case AC parameters by collecting data on a per-cycle basis from which the true results can be calculated.

To develop a more accurate measurement of SS output timings, cycle-to-cycle measurements need to be made. This is possible using the features of an ATE and post processing of data collected. For example, the tester’s capture memory and multiple passes through the pattern suite can be used to make a cycle-by-cycle comparison of the data to its strobe. The per-cycle comparison is identical to how the receiver in the actual system would operate, and it effectively removes all of the “common jitter” from the test. This results in a true measurement of the data to strobe skew including any cycle-by-cycle “differential jitter”.

Cycle-to-cycle measurement results are contrasted with the worst-case approach above in Figure 96. The graph illustrates the behavior of the Data and Strobe signals throughout the execution of one test pattern (part of the pattern is being displayed). It also shows the Tdvb parameter (the difference between these signals) as calculated by a post-processing script. The signals display a typical behavior of Source Synchronous signals; from cycle to cycle, both the data and the strobe signal delays change, maintaining approximately the same difference. A traditional search for a minimum value on Strobe would provide cycle A (5.15ns), while another traditional search for a maximum value on the Data would provide cycle B (3.6ns). From the graph, it can easily be understood that the difference between these values (1.55ns) is a meaningless Tdvb value, and is wrong to be considered worst-case. The data collected on a per-cycle basis, reveals that the worst-case Tdvb value is 1.73ns, which actually occurs at cycle C.

With modifications, the measurement technique described above can be used for HVM testing. One can first reduce the pattern set on which to take the measurements. Results from the design validation can be used to identify the worst cycles. Second, one can make the calculations on-line and more efficient by leveraging the ATE's capabilities. Finally, one can pursue a joint solution with an ATE vendor. The test and design guardband allowances in Tables 2-5 can be supported with this testing methodology.

3.1.3 Measuring Pad timings at a Pin

The AGP8X interface timings are specified at the pads. This choice was made to improve overall system timings. It presents another challenge to manufacturing tests as test equipment can only access the silicon via the pins. The solution is to translate the specs from the pad to the pin by accounting for the package effects. This section gives details for this solution, which consists of two parts. The first part translates between pad and pin timing specifications. The second part involves making changes to the tester calibration using the Test Interface Unit (TIU, a.k.a. loadboard).

Below, Figure 97 and Figure 98 illustrate the boundaries from silicon pad to package pin for driver and receiver specifications. With a straightforward application of the timing path, one can derive the equations to relate pad specifications to pin timings. These equations are listed below:

$$\begin{aligned} \text{Tsu @ pad} &= \text{Tsu @ pin} - \text{PackageR_Data_Longest_Delay} + \text{PackageR_GroupStrobe_Delay} \\ \text{Tho @ pad} &= \text{Tho @ pin} + \text{PackageR_Data_Shortest_Delay} - \text{PackageR_GroupStrobe_Delay} \\ \text{Tvb @ pad} &= \text{Tvb @ pin} - \text{PackageD_DataGroup_Longest_Delay} - \\ &\quad \text{PackageD_GroupStrobe_Delay} \\ \text{Tva @ pad} &= \text{Tva @ pin} - \text{PackageD_DataGroup_Shortest_Delay} + \\ &\quad \text{PackageD_GroupStrobe_Delay} \end{aligned}$$

With today's functional testers, any timing characterization or go-no go test is preceded by Time Domain Reflectometry (TDR) measurements taken in order to calibrate the tester for the delays not associated with the component. These delays are used to compensate the tester timings so that the tester signals are driven and component signals sensed with respect to time at the component socket.

Figure 99 depicts the TDR path for a single tester channel from tester pin electronics to the test socket. The path starts at the Tester Pin Electronics I/O. If there is no TIU (Test Interface Unit) attached to the test head, the TDR path ends at the pogo pin exposed on the test head. Since it is desirable to eliminate the TIU signal trace delays from the timing measurements along with the tester trace delays, the TIU is usually attached to the test head with a special vacuum seal before the TDR measurements are taken. The acquired delays from per-pin TDR measurements are stored to an EEPROM available on the TIU. The line delays are read off of the EEPROM at the initialization of the test program, and are used to calibrate the timing of the tester signal drive and strobe events, so that all timings are tested at the pins of the component plugged into this socket.

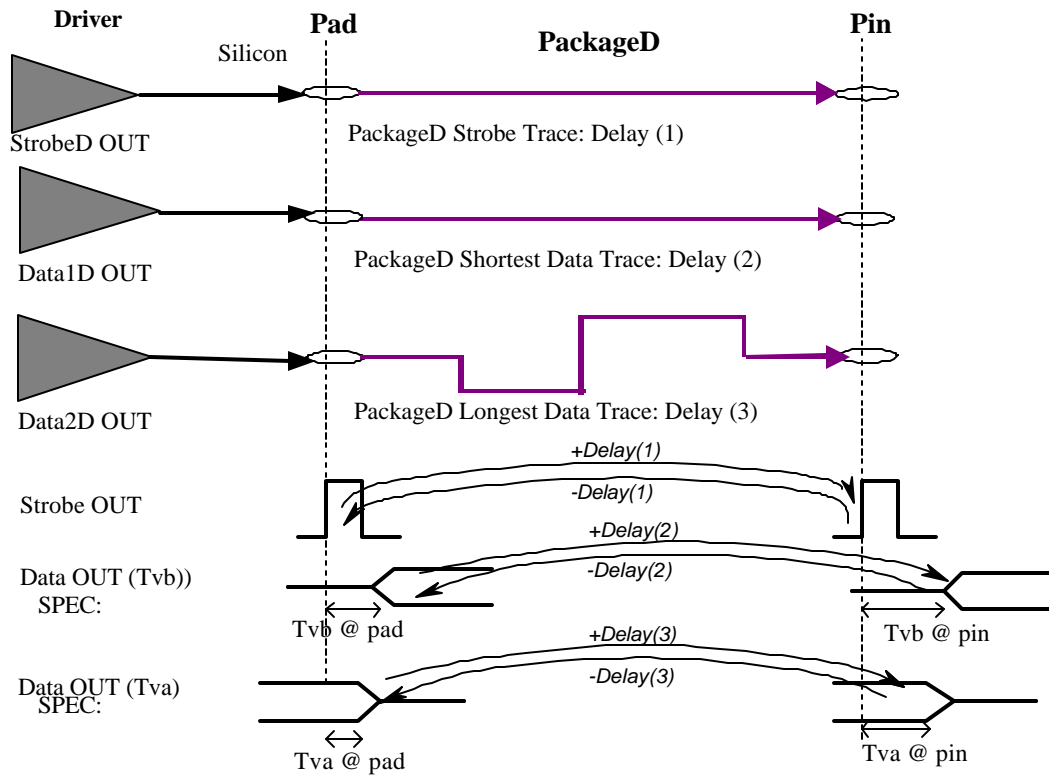


Figure 97: Pad to Pin Translations, Driver

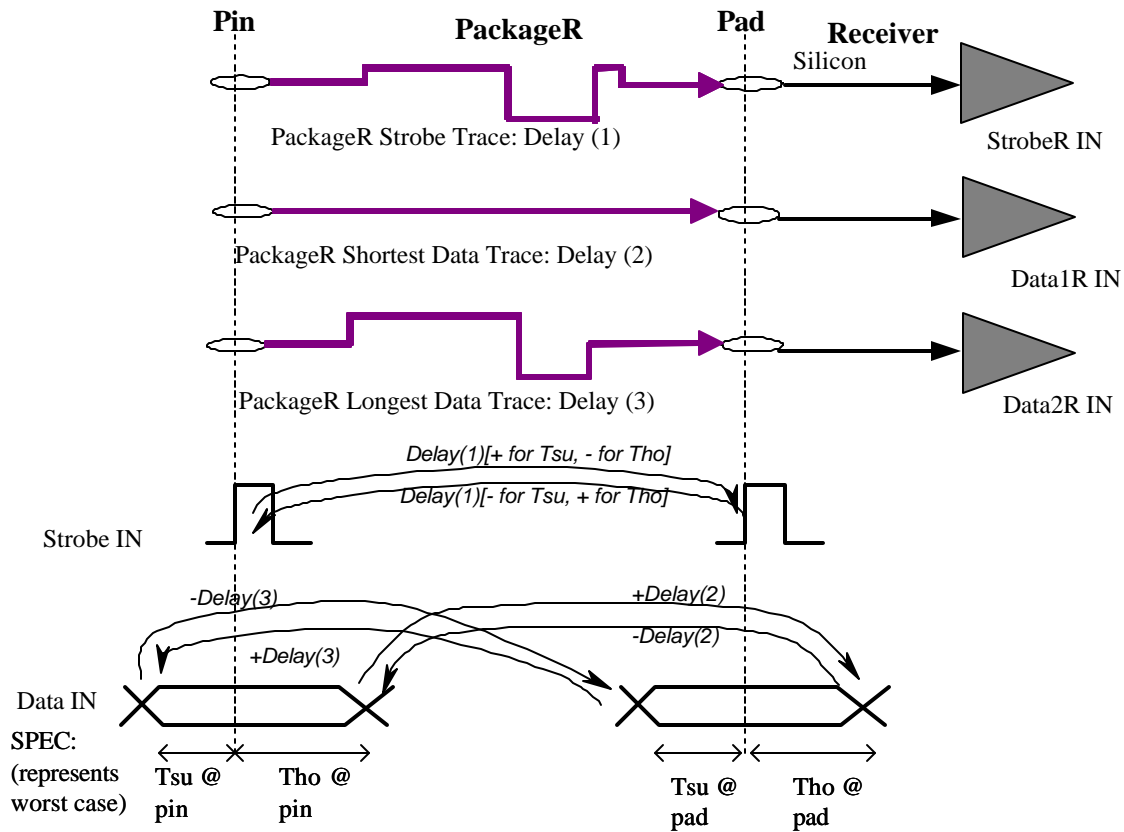


Figure 98: Pad to Pin Translations, Receiver

Based on the silicon timing specs at the pads, eliminating the package factor from the timing measurements is as simple as eliminating the timing delays associated with the package. This can be accomplished through plugging an open package (with no die inside) into the TIU test socket before the TDR routine is run. The final point of the TDR path then becomes the open end of the package trace as shown in Figure 100. The tester measurements thereafter compensate for time delays up to the die pad connected to this trace, as they would up to the package pin with the traditional component test method.

The test method heavily relies on the assumption that the open package socketed for timing calibrations is a good representative of the device packages socketed afterwards. In the first instance, this assumption should be true, since any major variation in package characteristics would result in the component timing specs to fail even with the current pin based methodology. This extra step of inserting an open package into the test socket before the regularly (weekly) executed TDR measurements does not have test time impact nor complicate the HVM test flow. The open package used for calibration can be updated when new revisions are available, as is the case with test program or silicon revisions.

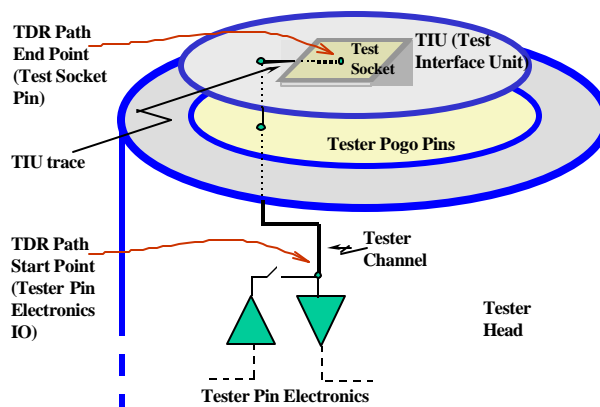


Figure 99: Conventional Tester Timing Calibrations

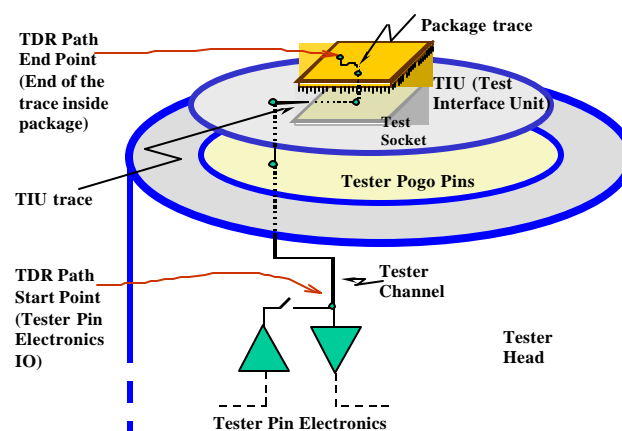


Figure 100: Tester Timing Calibrations for Pad Testing

For a wire-bond package, the bond wires may be small enough to ignore for the TDR process. If this is a concern, an alternative TDR method exists in which measurements are taken with all wirebond connections made to the ground plane in the package.

3.1.4 Testing with Low Data Rate Test Equipment

Your test equipment may not be able to provide the specified data rate of either the AGP8X or AGP4X interface. The HVM test goal test is to guarantee that your part will function at speed and that your part will meet all the timing requirements of the interface under test. A combination of a simple design for test (DFT) method and testing the timing specifications at a lower data rate can meet this test coverage requirement.

First, to verify that the interface functions at speed, DFT logic can be used to loopback the driver output to its own receiver input. A simple pattern can be used to check both by rising and falling transitions. Loopback is a fairly common DFT technique and one form of it is documented in a 1998 ITC paper by Gillis¹. Figure 101 below illustrates the scheme for a source synchronous interface.

¹ Gillis, P, et al, "Delay Test of Chip I/Os using LSSD Boundary," 1998 ITC Proceedings, pp83-90.

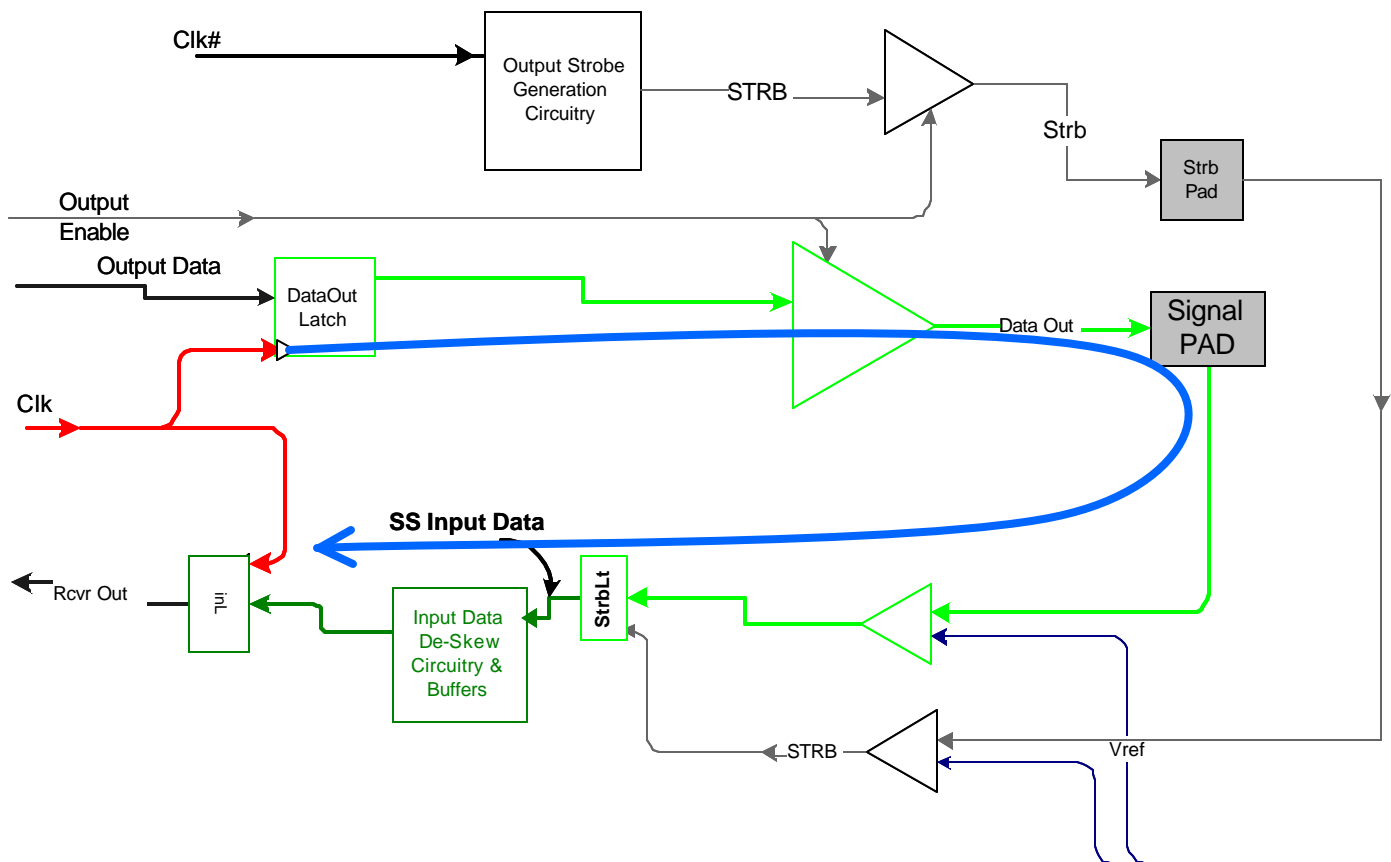


Figure 101: Loopback DFT for Source Synchronous Interface

The second component of guaranteeing timings on a low data rate tester is scaling the timing specifications to a data rate that is possible. For instance, consider testing the AGP8X interface at 266 MT/s instead of 533 MT/s. The input timings don't change, but the output timings will change because the period has doubled. Since the silicon doesn't change, the additional margin in the timings appears in the TDva and TDvb timings.

At 533 MT/s: TDva=527.5 ps and TDvb= 477.5 ps;

At 266 MT/s: TDva= 1465 ps and TDvb= 1415 ps.

The same test issues of EPA and pad specs discussed in previous sections about timings still hold true when evaluating timings at a lower data rate.

3.2 System Validation Measurement Guidelines

System validation requires different levels of verification starting at the electrical layer, progressing to the protocol layer and upwards. Validation of the electrical and protocol layers requires different types of measurements corresponding to a verification of voltage and timing limits for the electrical layer and the design of a logic analyzer interface (LAI) for the protocol layer and higher. This chapter is partitioned into a signal integrity section, which describes the equipment and measurement techniques to validate signal quality, and an LAI section, which describes the logic analyzer interface design requirements.

The next few sections explain how to measure voltage and timing margins at the receiver.

3.2.1 Measurement Techniques

Timing and voltage measurements should be made at the receiver pads, or as close to the receiver pads as possible.

3.2.2 Test Equipment Required

AGP8X is a high-speed parallel bus with a common clock speed of 66 MHz and a source synchronous rate of 533 MT/s. High performance test equipment must be used to obtain accurate data and to minimize loading of the circuit under test. Validating signal timings on AGP8X requires an oscilloscope with two or more channels with signal bandwidth of at least 1.5 GHz. Table 44 lists recommended measurement instruments.

Table 47: Recommended Time Domain Test Equipment

Mfg.	Type	Model	BW	C _L	R _L	Sample Rate	Notes
Tektronix	Oscilloscope	TDS694C	3 GHz	NA	NA	10 G	
Tektronix	Probes	P6249	4 GHz	1 pF	20 K Ω	NA	
Agilent Technologies	Oscilloscope	54846A	2.25 GHz	NA	NA	8.0 G	
Agilent Technologies	Probes	1152A	2.5 GHz	0.6 pF	100 K Ω	NA	

3.2.2.1 Oscilloscopes

Common clock measurements may be made with either a sampling mode or a real-time oscilloscope by connecting the common clock to the oscilloscope trigger and one input channel. The data is connected to the second channel. An eye diagram may be generated using the oscilloscope's infinite persistence mode and triggering repeatedly on the common clock.

For source synchronous timing measurements, a high bandwidth real-time oscilloscope is needed. To understand why, consider the case of data and strobe where the strobe's cycle-to-cycle edge placement may vary, but its skew with respect to the data remains very tight.

A sampling oscilloscope operates by acquiring data one point at a time over many cycles, slowly building up a composite image of these many samples. However, the data-to-strobe skew so acquired would not represent the data to strobe skew within a given bit period, but rather the worst-case skew between data of one bit cell and the strobe from another bit cell. Unlike a sampling oscilloscope, a real-time oscilloscope acquires a run of data at least as long as a single bit cell of data and then determines the data-to-clock skew. It then stores that information and acquires data for another bit cycle sometime later. By maintaining a database of several hundred to several thousand such samples, it is possible to generate a skew histogram.

3.2.2.2 Scope Probes

Low capacitance probes are essential. For best results, the probe capacitance should be 1.5 pF or less, and the DC impedance should be at least 10K ohms. Active probes best meet these constraints. Probe bandwidth should be at least 2.0 GHz. Active micro probes also work, since tip capacitances are

approximately 0.05 pF.

Probe grounding is critical. Most high bandwidth active oscilloscope probes have a coaxial probe tip where the ground conductor is formed by the outside cylinder of the coaxial cable. A short (< 500 mils) length of conductor can be used as a ground strap between the probe's ground conductor and the circuit board's ground plane. It is recommended that the probes be positioned and held in place by means of a probe holder. These holders permit accurate probe positioning and will maintain the probe in the desired position without having to hold it in place.

Probes and the input channels they are connected to also may exhibit slightly differing propagation delays. Before making critical measurements, especially on source synchronous data eyes, be sure to de-skew the probes and their associated input channels. If not properly de-skewed, the difference in propagation delays will introduce an error in timing skew measurements. Consult the oscilloscope manual for proper de-skewing methods.

3.2.2.3 De-Embedding Test Equipment Bandwidth

As edge rates approach the 0.5V/ns range, the finite bandwidth effects of the oscilloscope and probes may become large enough to yield misleading measurements. For example, a signal at a receiver may appear to have an insufficient rise time, when in fact the actual signal seen by the receiver is within spec. When a measurement is taken with an oscilloscope and oscilloscope probes, the displayed waveform reflects the effects of the finite bandwidth of the test equipment plus the actual signal's bandwidth. The effect of multiple bandwidth limiting stages may be expressed by the following equation:

$$1/t_{RO}^2 = 1/t_{scope}^2 + 1/t_{probe}^2 + 1/t_{signal}^2$$

Rearranging terms and solving for the signal rise time yields:

$$t_{signal} = [1/t_{RO}^2 - 1/t_{scope}^2 - 1/t_{probe}^2]^{-1/2}$$

where: t_{RO} = observed signal rise time

t_{scope} = oscilloscope rise time

t_{probe} = probe rise time

t_{signal} = actual signal rise time

Note that these equations do not take into account any reduction in edge rate caused by the capacitive loading by the probe. This is one reason why it is important to limit probe capacitive loading.

3.3 Time Domain Measurements

The following types of measurements should be made to obtain a complete view of the common clock and source synchronous timing and signal integrity.

3.3.1 66 MHz clock jitter and skew

The purpose of this measurement is to determine how much jitter and skew is present between the 66 MHz clock inputs of the two devices at each end of an AGP8X bus. If the clock jitter and skew measurements are too large, this points to a problem either with the 66 MHz clock generator or the clock distribution network. Meeting the jitter and skew targets is required to guarantee setup and hold times for the 66 MHz signals.

Common clock jitter and skew may be measured by connecting one channel and trigger to the 66 MHz clock pin of one AGP8X device and the other channel to the clock pin of the other AGP8X device. The oscilloscope trigger voltage should be set to V_{REF} , and the oscilloscope should be set in the infinite persistence mode. Many high performance oscilloscopes have an auto histogram mode that can be used to generate a clock jitter/skew histogram.

3.3.2 Common Clock Data Timing

In this test, the Equipment Under Test (EUT) should be configured to transmit in one direction only. One oscilloscope channel plus trigger is connected to the receiving device's common clock pin and the other channel to the data pin under test. The oscilloscope should be configured to trigger on the rising edge of the clock at a voltage equivalent to V_{REF} . Using the infinite persistence mode, an eye diagram is acquired

over many cycles. By measuring on the eye diagram the delay from trigger (clock's rising edge) to either side of the eye, one can obtain the worst-case setup and hold times.

3.3.3. Source Synchronous Timing

The best way of measuring source synchronous timing and voltage parameters is via a 2-channel oscilloscope capturing an eye diagram of the data while using the data strobe as a trigger source. Eye diagrams are particularly useful because they can be used to simultaneously display voltage and timing margins for a multiplicity of data patterns over time. The following diagram is an example of an eye diagram.





Figure 102: Eye Diagram Example with Strobe

For source synchronous timing, the EUT should be configured to transmit in one direction only. One oscilloscope channel plus trigger is connected to the receiving device's strobe pin and the other channel to the data pin under test. The oscilloscope should be configured to trigger on the rising edge of the clock at a voltage equivalent to V_{REF} . As described earlier, source synchronous signaling must be measured such that the data to strobe skew is acquired for an entire cycle.

The oscilloscope determines setup and hold times for source synchronous timing by filling its buffer with closely spaced samples that comprise several consecutive cell times. The sample rate is very high (8 GS/Sec or greater), and the buffer depth is relatively small (16 K samples). As a result, the oscilloscope is only able to buffer a relatively short burst of data. Once the data has been acquired, the oscilloscope's software interpolates to find the exact point where the data and strobe cross the specified threshold voltage, which should be preset to equal V_{REF} . The software then computes the T_{SU} and T_H within the sample interval and stores the values. The oscilloscope then acquires another interval of data and repeats the process, gradually producing a setup and hold time histogram. The oscilloscope setup procedure to measure data to strobe skew for a single cycle tends to be complex and may be found in the help menu of the particular oscilloscope in use.

Figure 103: Cycle-to-Cycle T_{SU} , T_H Measurement

Under the AGP8X signaling scheme, data is latched on the **rising** edge of the strobes. Because there are two strobes associated with a signal group, it is necessary to make separate measurements for data to AD_STBF and data to AD_STBS skew. Similarly, separate measurements must be taken for data to SB_STBF and data to SB_STBS skew.

3.4 Frequency Domain Measurements

When attempting to characterize the Power Distribution Network (PDN), it is easier to use frequency domain analysis to detect resonance peaks. The advantage of frequency domain analysis is its ability to identify in the frequency domain data pattern sensitivity, without the need to run a large number of bit patterns in the time domain, with the hope that the worst-case pattern has been tried. A vector network analyzer (VNA) is usually employed for making these measurements.

3.4.1 VNA Measurements

The best way of making measurements of the PDN is to probe directly across the V_{CC} and V_{SS} rails of the driver on the chip. If this measurement is not possible, it is possible to measure across the package's power and ground die attach points. Either measurement will require the use of micro-tip coplanar probes. Such probes are available from GGB Industries and may be ordered with different probe tip dimensions. The VNA and the probes must be calibrated prior to making a measurement. The VNA should be set to display an impedance vs. frequency plot.

3.4.2 Interpreting and Using the Data

S parameter plots display both real and imaginary components of a PDN impedance. The following figure is a sample plot for a PDN.




Figure 104: Sample S-parameter Plot for PDN

The plot displays impedance as a vector in which the distance above the x-axis specifies the imaginary part of the impedance, and the distance from the unit circle defines the real part. A point above the x-axis looks inductive, while a point below the x-axis looks capacitive. The distance radially inward from the larger unit circle defines the resistive component. In this example, the frequency was swept from 50 MHz (the low limit of the VNA) up to 2.0 GHz. The 2.0 GHz limit is near the limit beyond which there is very little spectral energy. At 50 MHz, the PDN looks almost purely capacitive with a very small resistive component. From 50 MHz to 500 MHz, the PDN remains capacitive, but experiences a slight increase in its resistive component. From 500 MHz to 1000 MHz, the network starts looking less capacitive while its resistance increases. Slightly above 1000 MHz, the PDN displays a resonance as indicated by the PDN changing from capacitive to inductive. As the frequency increases from 100 MHz to 2000 MHz, there is another resonance around 1700 MHz.

3.5 Waveform Measurements and Parameters

This section describes the proper method of extracting timing and voltage data from measurements. Because the scope of this document is system level validation, measurements will be specified at the **receiver**, thereby including the effects of the driver, interconnect, etc. For a description of the driver requirements, consult the AG3.0 Specification, Revision 1.0.

3.5.1 66 MHz Clock Jitter and Skew

66 MHz clock jitter and skew are measured by the oscilloscope in the infinite persistence mode and triggering on the rising edge of the clock at either the graphics device or the core logic device. In this example, the skew between the two 66 MHz clocks at either end of an AGP8X interface are being measured. The oscilloscope's trigger voltage should be set to the trigger voltage of the clock receiver. Clock jitter and skew are measured by comparing the spread that is seen at the second clock as shown in the diagram below.

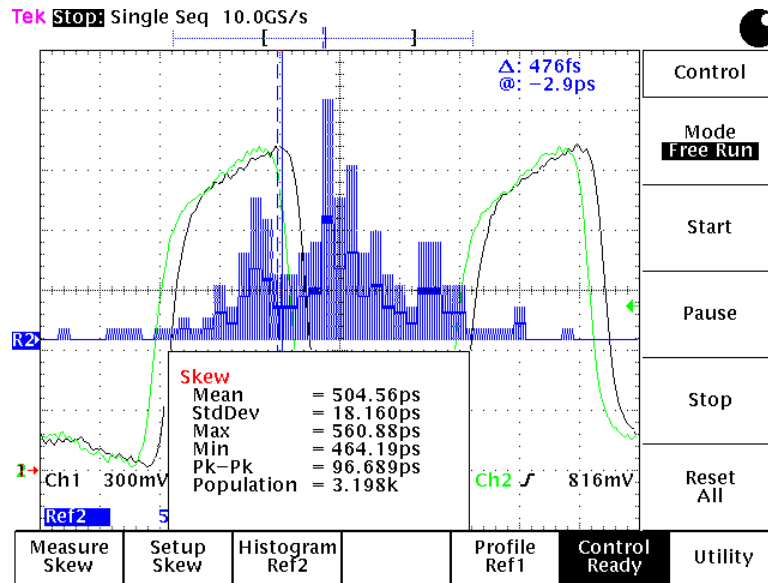


Figure 105: 66 MHz Clock Jitter and Skew Measurement Example

In this example, the histogram option is enabled, thereby giving more detailed information about the skew and jitter components. Skew correlates to the mean, while jitter is reflected in the standard deviation and peak-to-peak numbers.

3.5.2 Data to 66 MHz Clock Timing and Voltage Parameters

The following diagram shows how 66 MHz clock timing should be measured at the receiver. Setup and hold times are measured with the oscilloscope in the infinite persistence mode and the oscilloscope triggering on the rising edge of 66 MHz clock. Both data and clock are measured at the respective pins of the receiving device. The oscilloscope trigger should be set to the middle of the receiver's clock or to V_{REF} , and the data valid voltage should also be measured at V_{REF} .




Figure 106: 66 MHz Clock Timing at Receiver

3.5.3 Source Synchronous Timing and Voltage Parameters

AGP8X references its termination to V_{SS} . As a result, the bus idle state is at V_{SS} . The following diagram illustrates the waveform seen at the receiver for source synchronous data and strobes.




Figure 107: AGP8X Source Synchronous Data/strobe at Receiver

Note that all setup and hold times are referenced to the rising edge of a strobe. Since there are two strobes, it is necessary to make two data to strobe jitter/skew measurements: one for AD_STBF (SB_STBF) and the other for AD_STBS (SB_STBS). The oscilloscope trigger should be set to the middle of the receiver's clock or to V_{REF} , and the data valid voltage should also be measured at V_{REF} . As mentioned earlier, source synchronous measurements require a real-time oscilloscope with cycle-to-cycle jitter measurement capability.

3.6 Voltage Margin Measurements

These measurements made at the receiver should include the following

- Overshoot and undershoot limits
- Minimum Edge rate
- Ringback
- V_{IL} , V_{IH} limits

Acceptable limits for these parameters may be found in the AGP8X physical layer specifications. The method of measuring these parameters is the same as has been defined for measuring their simulated equivalents and may be found in section 2.3.2 of this document.

3.6.1 Data Test Patterns

Data test patterns should be selected to induce the worst-case timing and voltage margins. Data patterns that vary in frequency, such as 10101010, 110110110, 11001100, etc., can expose any data pattern induced resonance in the power distribution network. These patterns could be used to study the effects of Simultaneous Switching Outputs (SSO), Simultaneous Switching Inputs (SSI), Inter-Symbol Interference (ISI), even mode and odd mode cross talk, ground bounce, and voltage droop. The test software should try to drive these test patterns uninterrupted for the longest duration allowed by the test system. In general, the developer should produce patterns which are the same as those used for simulation purposes.

3.7 Logical Measurements

Logical measurements permit the bus transactions to be observed at the protocol level and above. When making logical measurements, it is already assumed that the electrical integrity of the bus has been validated. Due to the high signaling rate, it is necessary to design the Logic Analyzer Interface (LAI) into the system as part of the AGP8X system board layout process.

3.7.1 AGP8X Logic Analyzer Interface

Ideally, the LAI would exhibit no parasitic effects and would permit easy connection to the logic analyzer. In practice, the LAI is a compromise, offering a small loading on the bus and requiring some signal conditioning at the logic analyzer. One solution is to resistively isolate the LAI connector and regenerate the signals to the logic analyzer in the LAI probe. The value of the isolation resistors should be such that the current flowing through them is small compared to the current flowing through the AGP8X bus. Typically, a 10/1 ratio between the bus Z_0 and the isolation resistor is sufficient, so a value of 500 Ω is a

good first choice. The capacitive load of the resistor pad and the trace (and via if required) that connects the bus to the resistor must be minimized. The smaller the resistor package, the smaller the pads and hence the lower the parasitic capacitance. It is recommended that the resistors be physically no larger than the 0603 surface mount package.

The higher the bus speed, the greater sensitivity to bus loading. For example: the combined capacitive loading of the resistor pad and the trace for a 533 MT/s AGP8X data line should not exceed 1.0 pF, and the stub length should not exceed 200 mils. If the PCB technology supports it, via in pad and/or blind vias may be used to reduce the capacitance and stub length.

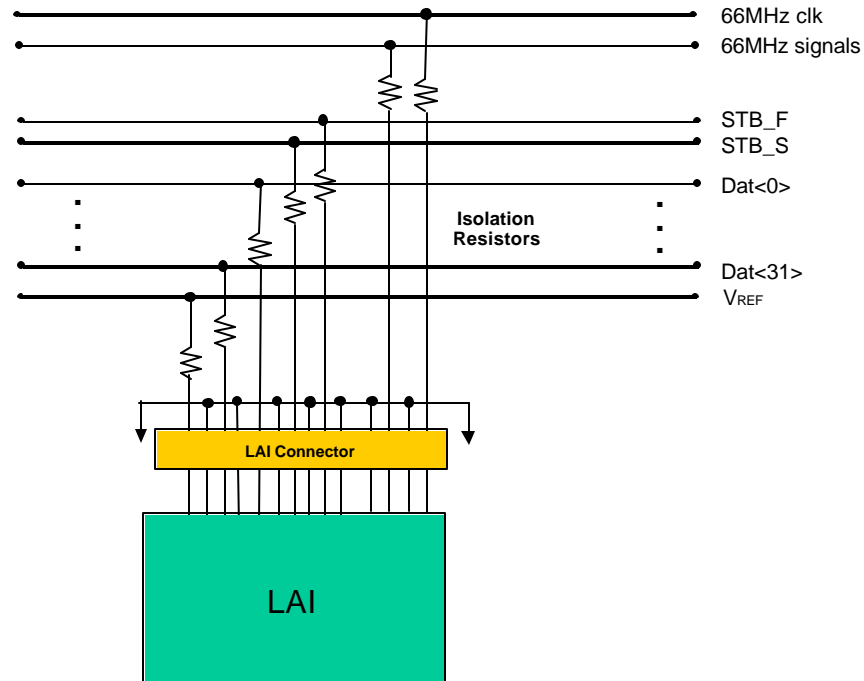


Figure 108: AGP8X Logic Analyzer Interface

The LAI attachment to the system board must be made in such a manner that no routing induced skew occurs. This requirement affects how source synchronous signals must be observed. They must be located on the bus such that the propagation delays between data and strobe are matched for both directions of data flow. This requirement is one of the reasons why it is necessary to separately match delays on AGP8X graphics cards as well as on motherboards. The following diagram illustrates one possible layout for implementing an LAI attachment.

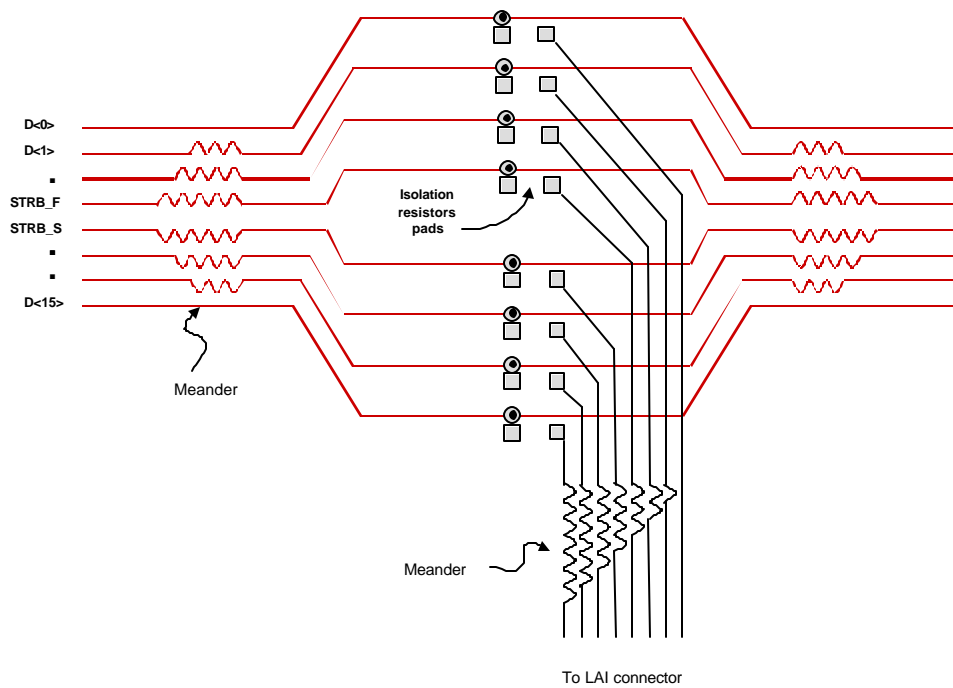


Figure 109: AGP8X Logic Analyzer Interface Connector Breakout

Common clock signals and the 66 MHz clock do not need to be delay matched. It is undesirable to add any additional delay on the common clock signals because this delay must be added to the common clock signals' t_{FLT} delay. Any impedance discontinuity on the common clock signals or the 66 MHz clock should be minimized because such discontinuities will also adversely affect t_{FLT} .

3.7.2 LAI Connector

The LAI connector must convey very low amplitude (~100 mVPP) signals to the LAI probe. Therefore the connector needs to maintain good signal integrity. It is recommended that each signal be surrounded on all sides by ground pins. Additionally, the overall connector length should be minimized.




Figure 110: LAI Keep-out Zone

The LAI connector and the isolation resistors may be directly placed on the system board. However, this is not always possible, particularly on production system boards. A special attachment mechanism may have to be designed that is compatible with the standard system board design. Since the AGP connector is the most likely place to attach the LAI, it is recommended that system board provide a 1.6 inch x 3 inch keep out zone around the AGP connector on the bottom side as shown in Figure 110. No components should be placed within this region on the bottom side of the system board. Contact the LAI vendors for details of the attachment method.

3.7.2.1 Recommended Probing Solution

A recommended method for probing the bottom side of the AGP connector is to use an anisotropic material as the conductive medium to provide a signal path from the AGP connector pin to an external test equipment input. Since the AGP connector does not provide for any type of mechanical attachment on the bottom side of the connector, a baseplate or housing could be permanently attached, using an adhesive, to the bottom side of the board surrounding the AGP connector. See Figure 111.




Figure 111: Probing Baseplate Housing Mounted Around Connector

The attached baseplate would then provide registration for a conductive copolymer interface material and attachment points for a interposer/interface board. See [Figure 112](#).




Figure 112: Baseplate Housing, Interposer and Anisotropic Material.

With a baseplate housing attached to the target board, and anisotropic material providing a conductive path, an interposer or interface board could be designed to mate with the baseplate housing. This interface board could contain signal buffering and logic required by the test equipment vendor as needed. See [Figure 113](#).

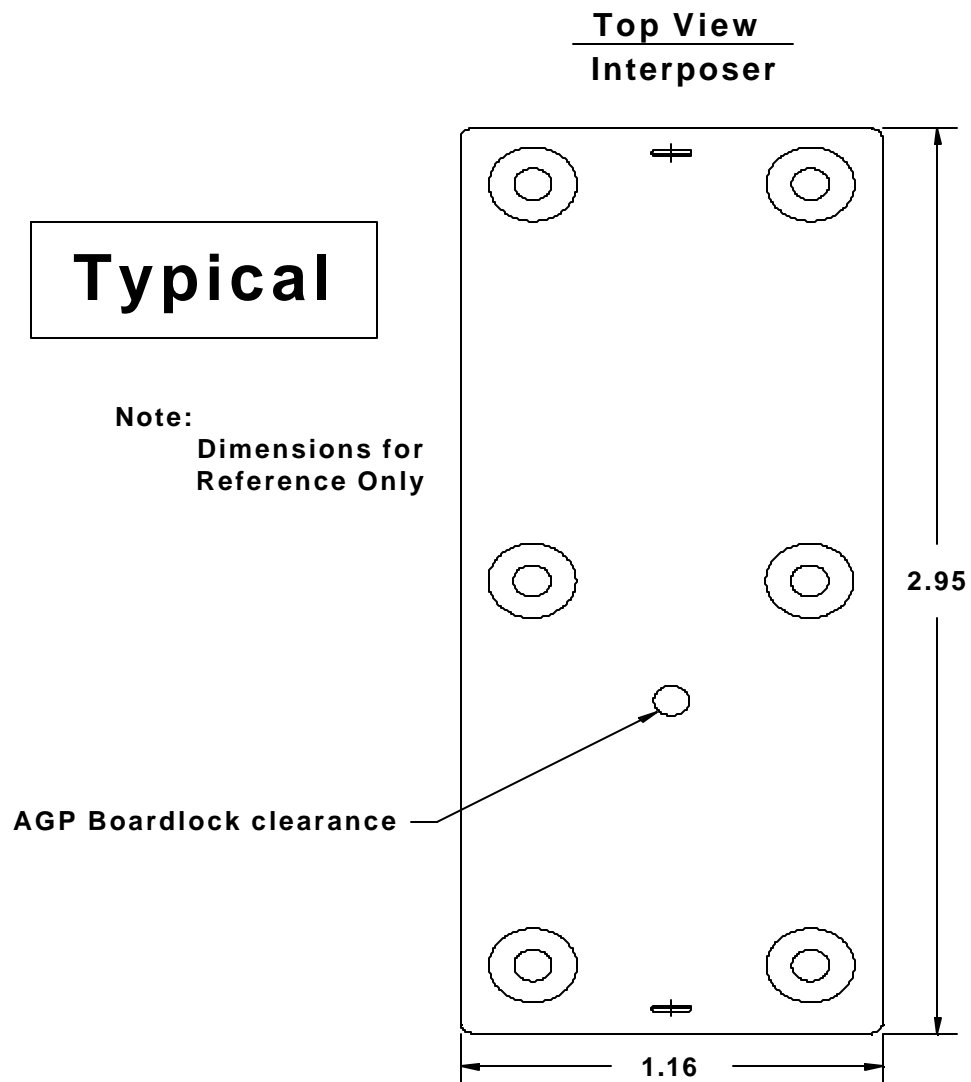


Figure 113: Interposer board.

4 AGP Thermal Design Guidelines

4.1 Thermal Overview

The general trend for high performance AGP-compliant graphics controllers is toward integrated and higher frequency functions; for example, 2-D, 3-D, and RAMDAC functions, in a single component. This trend will enable low component cost with high performance and enable flexibility for motherboard or add-in card solutions. For competitive 3-D performance levels, the graphics controller power trend will increase with time. The forecasted graphics controller maximum power levels at AGP introduction were 3-6 W. Future generation graphics controller power is expected to increase to 5-10 W for most controllers, with the potential of up to 15 W for some high performance controllers as shown in Figure 114

These power levels will create localized power densities that require thermal enhancement of component packages and add-in cards. Add-in card power levels have increased dramatically the last several years and in some cases may exceed the maximum of 25 watts as specified in the AGP Interface Specification, Revision 2.0. Add-in cards exceeding 25 watts must ensure adequate cooling is applied. Refer to section 4.2.




Figure 114: AGP Add-In Card and Graphics Controller Power Trend

4.2 Thermal Design Recommendations

4.2.1 Recommendations for Add-In Card Designers

It must be understood that standard thermal specifications such as ambient temperature and air flow cannot be controlled in certain market segments, in particular, the retail add-in card market segment. In this case, the worst case thermal environment expected is 50-60 °C internal ambient temperature near the AGP card with only natural convection air flow and any card orientation (vertical, horizontal-components up, horizontal-components down). For example, worst-case add-in card thermal performance typically occurs when the card orientation is horizontal with components facing down.

Add-in card component packages are likely to require thermal enhancement to operate in this type of

environment. The add-in card designer must determine the thermal environment of the target market segment and ensure that the add-in card is thermally compatible with this environment. In some cases this may include elements of active cooling.

4.2.2 Power Budgets

When designing a thermal solution for an add-in card, it is important to consider all major power consuming components on the entire graphics board. Considering that an average add-in graphics card can contain 4 MB of onboard local frame buffer (i.e., SGRAM, SDRAM, or DDR SDRAM), the power consumption of the memory section alone can exceed several watts. As the need for increased memory bandwidth arises in the 3D market segment, future memory technologies with higher operating frequencies and therefore higher power consumption will be available. In order to provide a complete thermal solution for an add-in graphics card, it is important to budget the total consumption of power on the add-in card, and then design an adequate cooling system based on the total number. This can be done by adding up the power consumption of the major components on the add-in card (i.e. graphics controller, memory, video, regulators, crystals, etc). As a result, the vendor should be able to design an adequate cooling system for the targeted market environment.

4.3 Recommendations for OEM System Designers

The OEM system integrator is strongly recommended to optimize the thermal environment for the AGP - compliant add-in card. This is typically accomplished by maximizing air flow or reducing ambient temperature. Improving the system thermal environment can reduce costs of add-in card thermal solutions which, in turn, may reduce overall system cost.

4.4 Thermal Performance Recommendations

4.4.1 Add-in Card




Figure 115: Airflow Over Add-In Card Components

If airflow is required over a densely populated add-in card, the inline orientation shown in Figure 115 is likely to produce higher ambient temperature for component B relative to component A. Component A blocks the airflow to component B in this configuration. In the staggered orientation shown in Figure 115, both components will equally benefit from the airflow.

4.4.2 Fan Direction

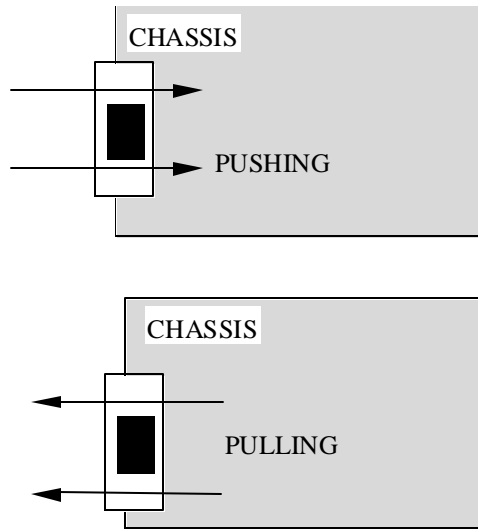


Figure 116: Pulling and Pushing Fan Airflow

Pushing Fan

- The pushing fan shown in Figure 116 can be incorporated in conjunction with a second fan of higher or lower CFM rating in both depressurized and pressurized chassis. It usually allows for a flexible device arrangement since it provides laminar airflow over a large area.
- Airflow is usually uniform and may eliminate most stagnant air and hot spots. It also keeps dust out. A front system fan is an example of a pushing fan. A pushing fan may slightly warm up the incoming air with the fan motor heat dissipation.

Pulling Fan

- The pulling fan shown in Figure 116 is generally used to exhaust heated air out of a chassis usually in conjunction with a pushing fan and, therefore, can be found in both depressurized and pressurized chassis. A power supply unit typically utilizes a pulling fan.
- Adequate venting is usually required for a pulling fan. The outgoing airflow is turbulent, which means better heat transfer in the turbulent region. However, this heat transfer area is limited around the pulling fan and may not be well defined throughout the system. Since this arrangement reduces pressure within a chassis, it tends to draw in dust through the vents and cracks in the chassis.

4.4.3 Venting




Figure 117: Adequate Venting Air

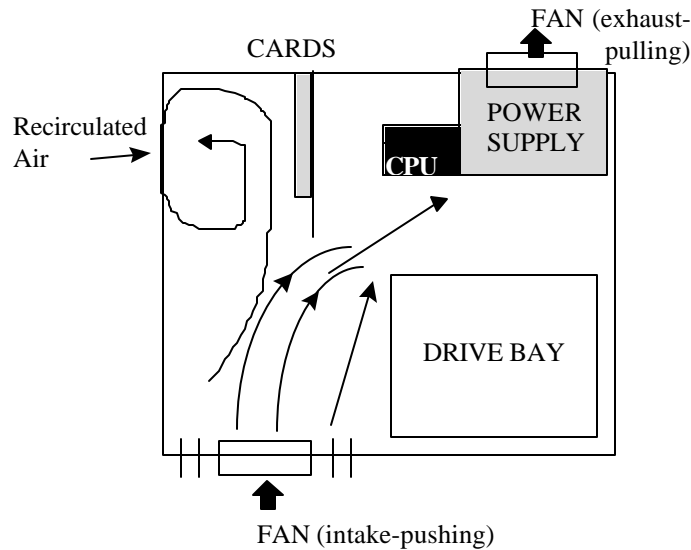


Figure 118: Inadequate Venting

Proper venting is a key element in any good thermal design. Proper venting will allow internal air to circulate and exchange with the ambient well as shown in Figure 117. Inadequate venting will cause internal air to recirculate as shown in Figure 118. If the chassis has been depressurized by fan action, vents placed at specific locations can permit an inflow of cooler air. This air can be made to flow in at certain critical locations.

Fan venting should be as generous as possible. Similarly, grille design should be as large as possible without conflicting with EMI and safety requirements.

4.4.4 Internal Air Flow

A full-length add-in card next to an AGP card is likely to block airflow to the graphics card as well as dissipate additional power. The AGP card cooling will be reduced by stagnant air and higher ambient temperatures. Therefore, it is recommended that a full-length card is not installed next to the AGP slot.

Cabling arrangement must not obstruct any component airflow or venting holes to ensure proper cooling.

4.5 Thermal Design Examples

The following examples of NLX and ATX chassis issues have been observed during thermal studies. Any other results will depend on the specific chassis, add-in card, and system configuration. The following examples are not representative of best or worst case thermal environments.

- **NLX Chassis Example:** Characterization of a representative NLX desktop form factor chassis thermal environment revealed low air flow (natural convection) and internal ambient temperatures of 60 °C near the AGP-enabled card. For this specific chassis form factor, the side panel vent shown in Figure 119 was added and resulted in a 20 °C drop in the internal ambient temperature near the AGP-enabled card. The side panel vent enabled the specific chassis design to cool a 20 W distributed thermal load card representing an AGP-enabled add-in card with little impact to the overall system thermal environment. The dimensions of the side panel vent are shown in Figure 120. An example of NLX rear vents is shown in Figure 121.




Figure 119: Example NLX Air Flow




Figure 120: Example NLX Side Panel Vent




Figure 121: Example NLX Rear Vents

- ATX Chassis Example:** Characterization of an ATX mini-tower chassis thermal environment revealed low air flow (natural convection) and internal ambient temperatures of 45 °C near a 12 W thermal load card representing an AGP add-in card. In this particular environment, component temperatures on the thermal load card were 20 °C above the upper specification limit. Addition of a front fan to the system, in addition to the existing power supply fan, increased the air flow in the vicinity of the AGP-enabled card and resulted in a drop of internal component temperatures to within the upper specification limit. The resulting air flow is shown in Figure 122.




Figure 122: Example ATX Air Flow

- Add-in Card Example:** A chassis design environment was analyzed where there was a 50 °C stagnant airflow condition surrounding a 12 W graphics card with a 5 W graphics controller in a horizontal card orientation with components down. The stagnant airflow condition resulted from adjacent add-in cards which blocked the natural convection air flow. Thermal simulation results indicated that a passive heat sink with dimensions of 3" x 2.3" x .5" attached to the graphics controller package met the controller thermal specification limits, but the maximum thermal specification limits of surrounding components such as cables (80 °C insulation limit), heat sink (U.L. 70 °C limit), and oscillator (data sheet 70 °C limit) were exceeded. For this specific thermal environment and power level, product safety and cost considerations dictated that a heat sink with active cooling would allow adequate controller cooling while meeting the thermal specifications of surrounding components.

4.5.1 Reference Documents

- AGP3.0 Interface Specification, Revision 1.0
- NLX Motherboard Specification
- ATX Specification
- Component Data Sheets and Application Notes

A. Crosstalk Modeling Techniques

A.1 Objective

In large multi-conductor systems, excessive line-to-line coupling, or crosstalk, can cause two detrimental effects. First of all, crosstalk will change the performance of the driven lines by modifying the effective characteristic impedance and propagation velocity, which will effect flight time and skew numbers. Additionally, crosstalk will induce noise onto other lines, which may degrade the signal integrity. The methodology presented in this appendix deals only with the first of these two effects. The study will determine how crosstalk can alter the skew and flight time of a transmission line. This appendix will also introduce simplified models that will reduce the simulation time of large N-conductor systems. Coupled lines in multi-conductor systems have variable impedances and propagation velocities which are dependent on coupling factors and data patterns. Because flight time skew is dependent on trace impedance and propagation velocities, crosstalk is an essential aspect of board design and analysis.

This section will explore the effects of odd and even propagation modes on flight time. It will be shown that simple single line models, with the appropriate parameters, can be used to model flight time differences due to line-to-line coupling in two-conductor systems. This concept will be expanded for use in the calculation of the effective impedance and propagation velocities in an N-conductor system. A methodology will be introduced that will allow an N-conductor system to be modeled as a single line for the purposes of flight time and skew analysis at greatly reduced simulation and analysis times.

It is important to note that this analysis is valid for linear systems only. A transmission line array can be approximated by a linear system only when frequencies are low. This technique should be valid for frequencies where the skin depth is larger than the conductor thickness. Above this frequency, non-linear effects such as internal inductance and frequency-dependent resistance will begin to be significant.

A.2 Flight Time and Skew as a Function of Odd and Even Propagation Modes

Simulations were performed in order to quantify signal skew and flight time as a function of propagation modes. The effective characteristic impedance and line delay will change with different propagation modes in systems where significant coupling between traces exist. Electric and magnetic fields interact in specific ways that are dependent on the data patterns on the coupled traces. These specific patterns can effectively reduce or increase the effective parasitic inductance and capacitance seen by a single line. Since the transmission line parameters vary with data patterns and are essential for accurate skew analysis, it is important to consider crosstalk analysis in high trace density/high speed systems. This appendix will introduce a skew/flight time analysis method where multiple conductors can be represented by a single conductor with appropriate variation in the impedance and velocity. This method should significantly reduce the simulation and analysis time needed to evaluate large systems.

A.2.1 Definitions

Characteristic Impedance and Propagation Delay: The Characteristic impedance and the propagation velocity are defined as follows:

$$\text{Characteristic impedance} = Z = \sqrt{\frac{L}{C}} \text{ ohms}$$

$$\text{Propagation delay (in seconds per unit distance)} = S = \sqrt{L C}$$

Where: L = inductance per unit length, and C = capacitance per unit length.

Transmission Line Parameters: The parameters used to describe the electrical equivalent of a transmission line. In this text, only the characteristic impedance (Z) and the propagation delay (S) are considered.

Inductance and Capacitance Matrix The matrix of parasitics reported by a field solver that represent the transmission line array effective inductance and capacitance values.

$$\text{Inductance Matrix} = \begin{bmatrix} L_{11} & L_{12} & \cdots & L_{1N} \\ L_{21} & L_{22} & & \\ \vdots & & \ddots & \\ L_{N1} & & & L_{NN} \end{bmatrix}$$

L_{NN} = Self inductance of line N

L_{MN} = Mutual inductance between lines M and N

$$\text{Capacitance Matrix} = \begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1N} \\ C_{21} & C_{22} & & \\ \vdots & & \ddots & \\ C_{N1} & & & C_{NN} \end{bmatrix}$$

C_{NN} = Total capacitance seen by line N = capacitance of conductor N to ground plus all the mutual capacitance to other lines.

C_{NM} = Mutual Capacitance between conductors N and M

For example, consider the following two-conductor system (Figure 123).




Figure 123: Simple Two-Conductor System Used to Explain the Parasitic Matrixes

$$\text{Capacitance matrix} = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}$$

C_{11} = Total capacitance for line 1 = Capacitance of line 1 to ground (C_s) plus all the mutual capacitance (i.e., for the two-conductor system, $C_{11} = C_{s1} + C_{12}$).

C_{12} = Mutual capacitance between line 1 and line 2.

$$\text{Inductance matrix} = \begin{bmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{bmatrix}$$

L_{11} = Self inductance for line 1

L_{12} = Mutual inductance between line 1 and line 2.

Odd Mode Propagation: Propagation mode when two coupled transmission lines are driven with equal magnitude and 180 degrees out of phase with one another. The fields arrange themselves such that there is a virtual ground plane between the two signals, which effectively doubles the mutual capacitance. Since the current in the lines is always traveling in opposite directions, the self inductance is reduced by the mutual inductance. The effective odd mode capacitance and inductance for a 2 line symmetrical system are:

$$C_{\text{odd}} = C_s + 2C_{12} = C_{11} + C_{12}$$

$$L_{\text{odd}} = L_{11} - L_{12}$$

$$Z_{\text{odd}} = \sqrt{\frac{L_{\text{odd}}}{C_{\text{odd}}}}$$

$$S_{\text{odd}} = \sqrt{L_{\text{odd}} C_{\text{odd}}} .$$

Even Mode Propagation: Propagation mode when two coupled transmission lines are driven with equal magnitude in phase with each other. The fields arrange themselves such that there is a virtual magnetic wall between the two conductors, which eliminates the effect of the mutual capacitance. The mutual inductance, however, is added to the self inductance because the current in both lines is always traveling in the same direction. The effective even mode capacitance and inductance for a two line symmetrical system are:

$$C_{\text{even}} = C_s = C_{11} - C_{12}$$

$$L_{\text{even}} = L_{11} + L_{12}$$

$$Z_{\text{even}} = \sqrt{\frac{L_{\text{even}}}{C_{\text{even}}}}$$

$$S_{\text{even}} = \sqrt{L_{\text{even}} C_{\text{even}}} .$$

Flight Time: Time measured between the $\frac{1}{2}(VCC)$ points of an unloaded driver and at the far end of the transmission line (at the receiver).

Skew: Difference in flight time between two signal nets.

A.2.2 Flight Time Analysis of Two Conductors Using a Single Conductor Model

Two coupled conductors can be modeled as one conductor by determining the effective odd mode or even mode impedance and propagation delay of the transmission line pair and substituting these parameters into a single line model.

Simulations were performed using XNS to simulate the effective odd mode and even mode flight time for a 2 inch transmission line pair at width/space ratios of 5/5, 5/10, and 5/15 mils. The relative dielectric constant was simulated as 4.0, and the dielectric thickness was set to 5.0 mils. The simulations were performed with loads of 0, 1, and 10 pF. Figure 124 depicts the simulation setup.

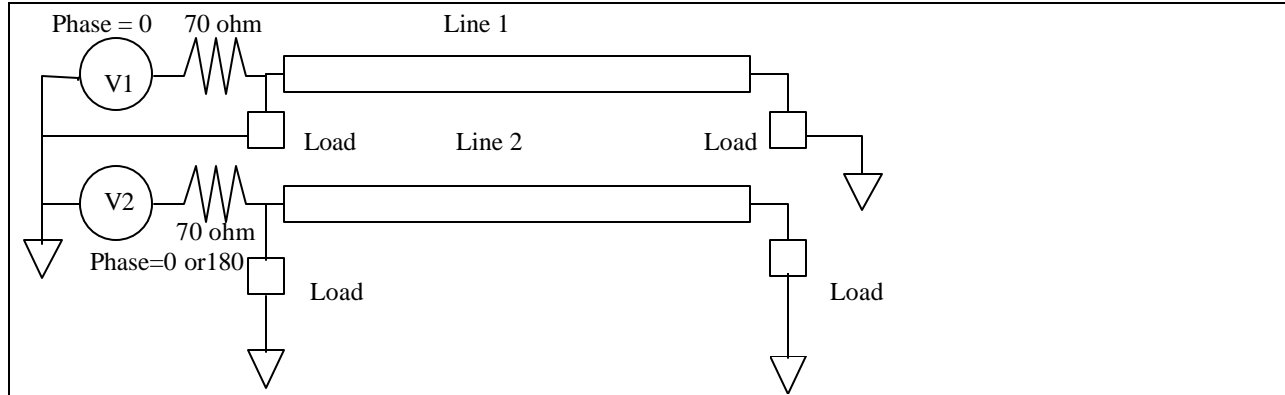


Figure 124: Schematic Used for Calculating Odd and Even Mode Flight Time Differences

XFX, which is a 2-D electromagnetic field simulator, was used to calculate the capacitance and inductance matrixes used in the simulation. These L and C matrixes were used to calculate the odd and even mode transmission line parameters using the above formulae. The resulting characteristic impedance and propagation delays were inserted into the single line model depicted in Figure 125. The flight time of the single line model was compared to the flight time of the coupled pair for both even and odd modes. Additionally, the loads were varied in the same manner as the coupled pair. The simulations confirm that single and differential pair simulation results match when the characteristics of the differential pair are reduced to single line equivalents using the equations in the above sections. The results are shown Table 48.

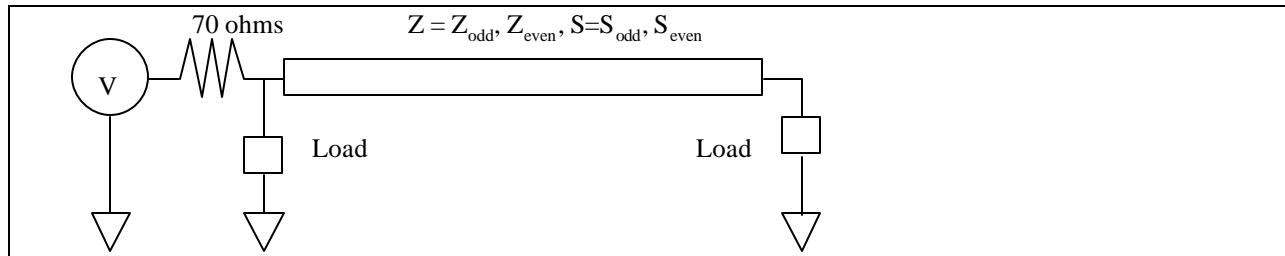


Figure 125: Schematic Used for Representing a Conductor Pair with a Single Line

Table 48: Results of Single Line and Coupled Line Simulations Showing Agreement Between Two-Conductor and One-Conductor Flight Times for Both Even and Odd Modes

Configuration	Simulated flight time in nanoseconds of odd and even propagation modes comparing coupled and single line models								
Spacing (mils)	5/5			5/10			5/15		
Load Capacitance	0 pF	1 pF	10 pF	0 pF	1 pF	10 pF	0 pF	1 pF	10 pF
Odd Mode Coupled Pair	0.27	0.30	0.52	0.27	0.30	0.54	0.28	0.31	0.56
Single Line with odd mode impedance and velocity	0.27	0.30	0.52	0.27	0.30	0.54	0.28	0.31	0.57
Even Mode Coupled Pair	0.29	0.33	0.63	0.29	0.32	0.61	0.29	0.32	0.6
Single Line with even mode impedance and velocity	0.29	0.33	0.63	0.29	0.32	0.61	0.29	0.32	0.6

The above results show that the single line equivalent model agrees with the two-conductor simulation. Furthermore, the simulations indicate that there can be significant differences in the flight times between a pair of signals that are propagating in phase and a pair of signals that are propagating 180 degrees out of phase. Additionally, when the load capacitance increases, the flight time differences between even and odd modes increases. It should also be noted that the flight time differences due to different propagation modes are exacerbated by tightly coupled lines.

A.2.3 Flight Time/Skew Analysis for More Than Two Conductors

The above simulations show how one conductor can replace two conductors for the purposes of incorporating crosstalk into flight time skew analysis. The objective of this study, however, is to determine a methodology to simulate one conductor in place of N conductors. To do so, formulae must be developed in order to correctly predict the equivalent inductance and capacitance for a given switching pattern of N bits on N coupled transmission lines.

A.2.4 Methodology

A methodology was developed to determine the effective impedance and velocity for multiple conductor lines. Initially, the equivalent capacitance and inductance seen by the target line (the line that you are doing the skew simulation on) must be determined for a specific switching pattern. This can be accomplished by observing pairs of lines in an N conductor system and determining the equivalent odd or even mode capacitance and inductance values. The values from the pairs can then be combined to produce the final value. To do so, the mutual components are either added or subtracted from the total capacitance or self inductance of the target line. When the signal line is switching in phase with the target line, then the mutual capacitance between the two lines is subtracted from the total capacitance of that line, and the mutual inductance is added. Conversely, when the signal line is switching out of phase with the target line, then the mutual capacitance between those two lines is added to the total capacitance of that line, and the mutual inductance is subtracted. The effective capacitance and inductance are used to calculate the equivalent characteristic impedance and propagation velocity for the single line equivalent model.

A.2.4.1 Example 1

Assume that switching pattern of the shown three conductor system is such that all bits on the net are switching in phase. Refer to the cross section in Figure 126.




Figure 126: Three Conductor Even Mode System

If the primary goal is to simulate this multi-conductor system with one line, and the line of interest is line 2, then the equivalent capacitance and inductance seen by line 2 can be calculated as follows:

Even Mode Capacitance of conductors 2 and 1 = $C_{s2} + C_{12} - C_{12} = C_{22} - C_{12}$

Even Mode Capacitance of conductors 2 and 3 = $C_{s2} + C_{13} - C_{13} = C_{22} - C_{23}$

Equivalent capacitance of middle conductor when conductors 1-3 are switching in phase = $C_{22} - C_{12} - C_{23}$

Even Mode Inductance of conductors 2 and 1 = $L_{22} + L_{12}$

Even Mode Inductance of conductors 2 and 3 = $L_{22} + L_{23}$

Equivalent inductance of middle conductor when conductors 1-3 are switching in phase = $L_{22} + L_{12} + L_{23}$.

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{2,eff} = \sqrt{\frac{L_{22} + L_{12} + L_{23}}{C_{22} - C_{12} - C_{23}}}$$

$$S_{2,eff} = \sqrt{(L_{22} + L_{12} + L_{23})(C_{22} - C_{12} - C_{23})}.$$

It is interesting to note that the mutual inductances are added or subtracted in the opposite manner as the mutual capacitance. This is due to the assumption that the system is linear and propagating only in TEM mode, which means that the electric and magnetic fields are always orthogonal to each other. As a result of the TEM assumption, the product of L and C remains constant for homogeneous systems. Thus, in a multi-conductor homogeneous system, such as a stripline array, if L is increased by the mutual inductance, then C must be decreased by the mutual capacitance such that LC remains constant.

In a non-homogeneous multi-conductor system, such as an array of microstrip lines, LC is not held constant for different propagation modes because the electromagnetic fields are traveling partially in air and partially in the dielectric material of the board. In a microstrip system, the effective dielectric constant is a weighted average between air and the dielectric material of the board. Because the field patterns change with different propagation modes, the effective dielectric constant will change depending on the field densities contained within the board dielectric material and the air. Thus, the LC product will be mode dependent in a non-homogeneous system. The LC product, however, will remain constant for a given mode.

A.2.4.2 Example 2

Assume that the switching pattern of a three conductor system is such that bits 1 and 2 are switching in phase and bit 3 is switching 180 degrees out of phase. (Line 2 is still the line of interest.) Figure 127 shows this system.

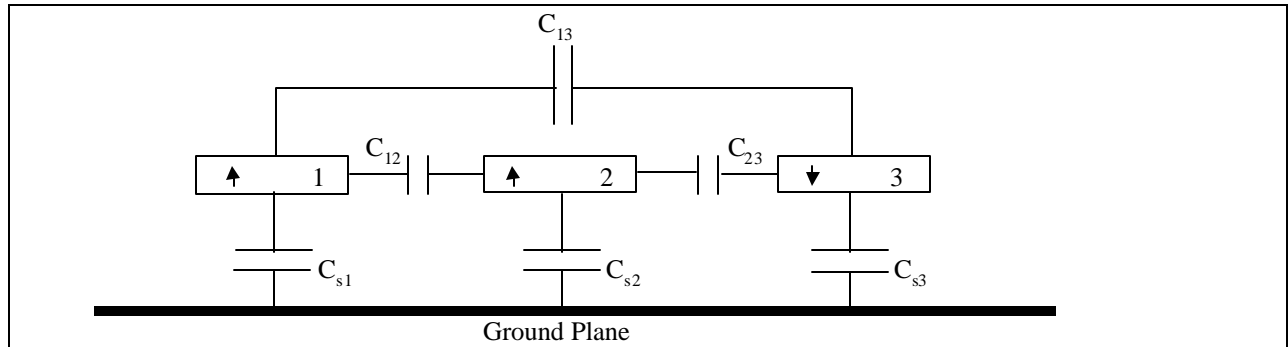


Figure 127: Three Conductor System with Odd Mode and Even Mode

Even Mode Capacitance of conductors 2 and 1 = $C_{22} - C_{12}$

Odd Mode Capacitance of conductors 2 and 3 = $C_{22} + C_{23}$

Equivalent capacitance of conductor 2 = $C_{22} - C_{12} + C_{23}$

Even Mode Inductance of conductors 2 and 1 = $L_{22} + L_{12}$

Odd Mode Inductance of conductors 2 and 3 = $L_{22} - L_{23}$

Equivalent inductance of conductor 2 = $L_{22} + L_{12} - L_{23}$.

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{2,eff} = \sqrt{\frac{L_{22} + L_{12} - L_{23}}{C_{22} - C_{12} + C_{23}}}$$

$$S_{2,eff} = \sqrt{(L_{22} + L_{12} - L_{23})(C_{22} - C_{12} + C_{23})}.$$

A.2.4.3 Example 3

Assume that switching pattern of a five conductor system is such that bits 1, 2, and 5 are switching in phase, and bits 3 and 4 are in phase with each other but 180 degrees out of phase with bits 1, 2, and 5. (Line 3 is the line of interest.) Figure 128 shows this system.




Figure 128: Five Conductor System

Odd Mode Capacitance of conductors 3 and 1 = $C_{33} + C_{13}$

Odd Mode Capacitance of conductors 3 and 2 = $C_{33} + C_{23}$

Even mode capacitance of conductors 3 and 4 = $C_{33} - C_{34}$

Odd mode capacitance of conductors 3 and 5 = $C_{33} + C_{35}$

Equivalent capacitance of conductor 3 = $C_{33} + C_{13} + C_{23} - C_{34} + C_{35}$

Odd Mode inductance of conductors 3 and 1 = $L_{33} - L_{13}$

Odd Mode inductance of conductors 3 and 2 = $L_{33} - L_{23}$

Even Mode inductance of conductors 3 and 4 = $L_{33} + L_{34}$

Odd mode inductance of conductors 3 and 5 = $L_{33} - L_{35}$

Equivalent inductance of conductor 3 = $L_{33} - L_{13} - L_{23} + L_{34} - L_{35}$.

The equivalent single conductor model that could be used to perform skew analysis of this multi-conductor system would have the following transmission line parameters:

$$Z_{3,eff} = \sqrt{\frac{L_{33} - L_{13} - L_{23} + L_{34} - L_{35}}{C_{33} + C_{13} + C_{23} - C_{34} + C_{35}}}$$

$$S_{3,eff} = \sqrt{(L_{33} - L_{13} - L_{23} + L_{34} - L_{35})(C_{33} + C_{13} + C_{23} - C_{34} + C_{35})}.$$

Simulations were performed on a three-conductor system to confirm that the above methodology produces the correct flight time results when a single line is used to represent a multi-conductor system. Table 49 shows simulation results comparing a three-conductor flight time simulation and a single conductor flight time simulation where the equivalent impedance and propagation delay have been used for the appropriate switching pattern. The target conductor is line 2 (L2). The simulations were performed using XNS to simulate the effective odd mode and even mode flight times for an array of 2-inch long transmission lines at width/space ratios of 5/5 mils. The relative dielectric constant was simulated as 4.0, and the dielectric thickness was set to 5.0 mils. The simulations were performed with loads of 0, 1, and 10 pF. The lines were driven with 1.8 volts and a 100 ps edge rate. See Figure 15 and Table 7 or Table 8 for a representative board cross section.

Table 49: Results of Multi-line and Single Line Flight Time Simulations Showing Agreement Between 3-Conductor and Equivalent Single Conductor Models

Switching Pattern L1 L2 L3	Effective Parasitics		Load = 0pF		Load = 1pF		Load = 10pF	
	L	C	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)	Flight Time (L2 of Multi-Line)	Flight Time (Equiv. Single Line)
1 0 0	L22-L12+ L23	C22+C12- C23	0.28 ns	0.28 ns	0.31 ns	0.31 ns	0.57 ns	0.57 ns
1 1 0	L22+L12- L23	C22-C12+ C23	0.28 ns	0.28 ns	0.31 ns	0.31 ns	0.57 ns	0.57 ns
1 0 1	L22-L12- L23	C22+C12+ C23	0.25 ns	0.25 ns	0.26 ns	0.26 ns	0.46 ns	0.46 ns
1 1 1	L22+L12+ L23	C22-C12- C23	0.30 ns	0.3 ns	0.34 ns	0.34 ns	0.67 ns	0.67 ns

B. Models

Table 50: AGP-4X Buffer I-V Curves

48 ohm curvature 2.0		38 ohm curvature 2.0		30 ohm curvature 1.1	
Voltage (V)	Current (mA)	Voltage (V)	Current (mA)	Voltage (V)	Current (mA)
0.0	0.0	0.0	0.0	0.0	0.0
0.05	1.424	0.073	3.122	0.132	4.94
0.1	2.811	0.144	5.969	0.256	9.446
0.15	4.157	0.212	8.54	0.373	13.516
0.2	5.462	0.279	10.837	0.483	17.15
0.25	6.72	0.344	12.86	0.586	20.351
0.3	7.928	0.407	14.611	0.682	23.123
0.35	9.082	0.467	16.094	0.77	25.47
0.4	10.178	0.526	17.313	0.852	27.399
0.45	11.209	0.583	18.273	0.927	28.918
0.5	12.171	0.639	18.98	0.996	30.038
0.55	13.056	0.692	19.443	1.058	30.77
0.6	13.857	0.744	19.732	1.115	31.228
0.65	14.564	0.796	19.986	1.172	31.629
0.7	15.166	0.848	20.21	1.228	31.984
0.75	15.652	0.899	20.407	1.283	32.296
0.8	16.008	0.95	20.58	1.337	32.57
0.85	16.241	1.001	20.731	1.391	32.808
0.9	16.443	1.052	20.863	1.445	33.017
0.95	16.625	1.103	20.978	1.498	33.198
1.0	16.787	1.154	21.079	1.550	33.358
1.05	16.931	1.205	21.168	1.603	33.5
1.1	17.058	1.255	21.249	1.655	33.628
1.15	17.17	1.306	21.324	1.707	33.747
1.2	17.269	1.356	21.396	1.759	33.86
1.25	17.356	1.407	21.466	1.811	33.971

AGP Design Guide, Rev. 2.0

1.3	17.434	1.457	21.535	1.862	34.08
1.35	17.503	1.508	21.603	1.914	34.188
1.4	17.567	1.558	21.67	1.966	34.294
1.45	17.626	1.609	21.735	2.018	34.398
1.5	17.684	1.659	21.8	2.069	34.5

Table 51: AGP-4X Diode Clamp I-V Curves

Voltage (V)	Current (mA)
0	0
1.5	0
2.2	0
2.4	10
2.6	100

AGP Design Guide, Rev. 2.0

The AGP-8X spec IBIS model pullup and pulldown curves are included here as a reference. The IBIS model is also plotted against the V-I curve spec described in *AGP3.0 Specification, Revision 1.0*. Notice the min/max refers to the buffer impedance. Voltage is in volts and current in Amperes.

Table 52: AGP 8X Buffer V-I Curves

[Pulldown]			
Voltage	I(typ)	I(min)	I(max)
-1.45E+00	-0.033144	0.000000	-0.029508
-1.40E+00	-0.030935	-0.032357	-0.027540
-1.30E+00	-0.028725	-0.030046	-0.025573
-1.20E+00	-0.026516	-0.027735	-0.023606
-1.10E+00	-0.024306	-0.025424	-0.021639
-1.00E+00	-0.022096	-0.023113	-0.019672
-9.00E-01	-0.019887	-0.020801	-0.017705
-8.00E-01	-0.017677	-0.018490	-0.015737
-7.00E-01	-0.015467	-0.016179	-0.013770
-6.00E-01	-0.013258	-0.013867	-0.011803
-5.00E-01	-0.011048	-0.011556	-0.009836
-4.00E-01	-0.008839	-0.009244	-0.007869
-3.00E-01	-0.006629	-0.006934	-0.005902
-2.00E-01	-0.004419	-0.004623	-0.003934
-1.00E-01	-0.002210	-0.002311	-0.001967
0.00E+00	0.000000	0.000000	0.000000
5.00E-02	0.001105	0.001156	0.000984
1.00E-01	0.002178	0.002288	0.001941
1.50E-01	0.003277	0.003480	0.002926
2.00E-01	0.004343	0.004645	0.003882
2.50E-01	0.005376	0.005785	0.004861
3.00E-01	0.006428	0.006898	0.005810
3.50E-01	0.007445	0.008062	0.006728
4.00E-01	0.008476	0.009123	0.007664
4.50E-01	0.009470	0.010304	0.008567
5.00E-01	0.010429	0.011384	0.009440
5.50E-01	0.011396	0.012504	0.010324
6.00E-01	0.012366	0.013527	0.011215
6.50E-01	0.013296	0.014654	0.012071
7.00E-01	0.014226	0.015684	0.012893
7.50E-01	0.015114	0.016747	0.013717
8.00E-01	0.015962	0.017778	0.014506
8.50E-01	0.016836	0.018835	0.015324
9.00E-01	0.017666	0.019858	0.016072
9.50E-01	0.018481	0.020900	0.016842
1.00E+00	0.019251	0.021854	0.017572
1.05E+00	0.020003	0.022875	0.018262
1.10E+00	0.020733	0.023809	0.018963
1.15E+00	0.021438	0.024754	0.019620
1.20E+00	0.022092	0.025706	0.020255
1.25E+00	0.022746	0.026658	0.020891
1.30E+00	0.023400	0.027609	0.021527
1.35E+00	0.024054	0.028561	0.022163
1.40E+00	0.024708	0.029513	0.022799
1.45E+00	0.025362	0.030465	0.023435

AGP Design Guide, Rev. 2.0

1.50E+00	0.026016	0.031416	0.024071
1.55E+00	0.026670	0.032369	0.024707
1.60E+00	0.027324	0.033320	0.025343
1.65E+00	0.027978	0.034273	0.025979
1.70E+00	0.028632	0.035224	0.026615
1.75E+00	0.029286	0.036176	0.027251
1.80E+00	0.029940	0.037128	0.027887
1.85E+00	0.030594	0.038080	0.028523
1.90E+00	0.031248	0.039032	0.029158
1.95E+00	0.031902	0.039984	0.029794
2.00E+00	0.032556	0.040935	0.030430
2.05E+00	0.033210	0.041888	0.031066
2.10E+00	0.033864	0.042839	0.031702
2.15E+00	0.034518	0.043792	0.032338
2.20E+00	0.035172	0.044743	0.032974
2.25E+00	0.035826	0.045695	0.033610
2.30E+00	0.036480	0.046647	0.034246
2.35E+00	0.037134	0.047599	0.034882
2.40E+00	0.037789	0.048551	0.035518
2.45E+00	0.038443	0.049503	0.036154
2.50E+00	0.039097	0.050454	0.036790
2.55E+00	0.039751	0.051407	0.037426
2.60E+00	0.040405	0.052358	0.038061
2.65E+00	0.041059	0.053311	0.038697
2.70E+00	0.041713	0.054262	0.039333
2.75E+00	0.042367	0.055214	0.039969
2.80E+00	0.043021	0.056166	0.040605
2.85E+00	0.043675	0.057118	0.041241
2.90E+00	0.044329	0.058069	0.041877
2.95E+00	0.044983	0.059022	0.042513
3.00E+00	0.045637	0.059973	0.043149

[PullUp]			
Voltage	I(typ)	I(min)	I(max)
-1.4500E+00	4.0421E-02	4.4795E-02	3.5317E-02
-1.4000E+00	3.9027E-02	4.3250E-02	3.4100E-02
-1.3500E+00	3.7633E-02	4.1705E-02	3.2881E-02
-1.3000E+00	3.6239E-02	4.0161E-02	3.1663E-02
-1.2500E+00	3.4846E-02	3.8616E-02	3.0446E-02
-1.2000E+00	3.3452E-02	3.7072E-02	2.9228E-02
-1.1500E+00	3.2058E-02	3.5527E-02	2.8011E-02
-1.1000E+00	3.0664E-02	3.3982E-02	2.6792E-02
-1.0500E+00	2.9270E-02	3.2437E-02	2.5574E-02
-1.0000E+00	2.7876E-02	3.0893E-02	2.4357E-02
-9.5000E-01	2.6483E-02	2.9348E-02	2.3139E-02
-9.0000E-01	2.5089E-02	2.7804E-02	2.1921E-02
-8.5000E-01	2.3695E-02	2.6259E-02	2.0703E-02
-8.0000E-01	2.2301E-02	2.4715E-02	1.9485E-02
-7.5000E-01	2.0907E-02	2.3170E-02	1.8268E-02
-7.0000E-01	1.9514E-02	2.1626E-02	1.7050E-02
-6.5000E-01	1.8120E-02	2.0080E-02	1.5832E-02
-6.0000E-01	1.6726E-02	1.8535E-02	1.4614E-02
-5.5000E-01	1.5332E-02	1.6991E-02	1.3396E-02

AGP Design Guide, Rev. 2.0

-5.0000E-01	1.3938E-02	1.5446E-02	1.2178E-02
-4.5000E-01	1.2544E-02	1.3902E-02	1.0961E-02
-4.0000E-01	1.1151E-02	1.2357E-02	9.7426E-03
-3.5000E-01	9.7570E-03	1.0813E-02	8.5244E-03
-3.0000E-01	8.3630E-03	9.2682E-03	7.3072E-03
-2.5000E-01	6.9690E-03	7.7237E-03	6.0890E-03
-2.0000E-01	5.5750E-03	6.1781E-03	4.8718E-03
-1.5000E-01	4.1810E-03	4.6336E-03	3.6536E-03
-1.0000E-01	2.7880E-03	3.0891E-03	2.4354E-03
-5.0000E-02	1.3940E-03	1.5445E-03	1.2182E-03
0.0000E+00	0.0000E+00	0.0000E+00	0.0000E+00
5.0000E-02	-1.3940E-03	-1.5445E-03	-1.2182E-03
1.0000E-01	-2.7450E-03	-3.0783E-03	-2.4257E-03
1.5000E-01	-4.0550E-03	-4.5649E-03	-3.5954E-03
2.0000E-01	-5.3510E-03	-6.0033E-03	-4.7254E-03
2.5000E-01	-6.5740E-03	-7.4245E-03	-5.8167E-03
3.0000E-01	-7.7820E-03	-8.7652E-03	-6.8934E-03
3.5000E-01	-8.9700E-03	-1.0086E-02	-7.9274E-03
4.0000E-01	-1.0087E-02	-1.1356E-02	-8.8995E-03
4.5000E-01	-1.1182E-02	-1.2602E-02	-9.8734E-03
5.0000E-01	-1.2207E-02	-1.3768E-02	-1.0784E-02
5.5000E-01	-1.3208E-02	-1.4907E-02	-1.1654E-02
6.0000E-01	-1.4182E-02	-1.5991E-02	-1.2500E-02
6.5000E-01	-1.5087E-02	-1.7022E-02	-1.3301E-02
7.0000E-01	-1.5962E-02	-1.8020E-02	-1.4059E-02
7.5000E-01	-1.6786E-02	-1.8940E-02	-1.4772E-02
8.0000E-01	-1.7560E-02	-1.9824E-02	-1.5456E-02
8.5000E-01	-1.8283E-02	-2.0651E-02	-1.6080E-02
9.0000E-01	-1.8970E-02	-2.1436E-02	-1.6672E-02
9.5000E-01	-1.9590E-02	-2.2146E-02	-1.7229E-02
1.0000E+00	-2.0184E-02	-2.2812E-02	-1.7727E-02
1.0500E+00	-2.0712E-02	-2.3418E-02	-1.8189E-02
1.1000E+00	-2.1197E-02	-2.3977E-02	-1.8603E-02
1.1500E+00	-2.1620E-02	-2.4464E-02	-1.8971E-02
1.2000E+00	-2.2001E-02	-2.4903E-02	-1.9293E-02
1.2500E+00	-2.2341E-02	-2.5295E-02	-1.9575E-02
1.3000E+00	-2.2629E-02	-2.5627E-02	-1.9825E-02
1.3500E+00	-2.2887E-02	-2.5924E-02	-2.0043E-02
1.4000E+00	-2.3116E-02	-2.6182E-02	-2.0239E-02
1.4500E+00	-2.3324E-02	-2.6421E-02	-2.0417E-02
1.5000E+00	-2.3517E-02	-2.6642E-02	-2.0584E-02
1.5500E+00	-2.3700E-02	-2.6851E-02	-2.0742E-02
1.6000E+00	-2.3874E-02	-2.7051E-02	-2.0896E-02
1.6500E+00	-2.4043E-02	-2.7243E-02	-2.1043E-02
1.7000E+00	-2.4207E-02	-2.7433E-02	-2.1185E-02
1.7500E+00	-2.4369E-02	-2.7614E-02	-2.1323E-02
1.8000E+00	-2.4524E-02	-2.7791E-02	-2.1460E-02
1.8500E+00	-2.4680E-02	-2.7968E-02	-2.1599E-02
1.9000E+00	-2.4835E-02	-2.8145E-02	-2.1736E-02
1.9500E+00	-2.4991E-02	-2.8321E-02	-2.1874E-02
2.0000E+00	-2.5146E-02	-2.8498E-02	-2.2012E-02
2.0500E+00	-2.5301E-02	-2.8675E-02	-2.2150E-02
2.1000E+00	-2.5457E-02	-2.8852E-02	-2.2288E-02

2.1500E+00	-2.5612E-02	-2.9029E-02	-2.2426E-02
2.2000E+00	-2.5768E-02	-2.9206E-02	-2.2563E-02
2.2500E+00	-2.5923E-02	-2.9383E-02	-2.2702E-02
2.3000E+00	-2.6079E-02	-2.9559E-02	-2.2839E-02
2.3500E+00	-2.6234E-02	-2.9736E-02	-2.2977E-02
2.4000E+00	-2.6389E-02	-2.9912E-02	-2.3115E-02
2.4500E+00	-2.6545E-02	-3.0089E-02	-2.3253E-02
2.5000E+00	-2.6700E-02	-3.0266E-02	-2.3391E-02
2.5500E+00	-2.6856E-02	-3.0443E-02	-2.3528E-02
2.6000E+00	-2.7011E-02	-3.0620E-02	-2.3666E-02
2.6500E+00	-2.7166E-02	-3.0797E-02	-2.3805E-02
2.7000E+00	-2.7322E-02	-3.0974E-02	-2.3942E-02
2.7500E+00	-2.7477E-02	-3.1151E-02	-2.4080E-02
2.8000E+00	-2.7633E-02	-3.1327E-02	-2.4217E-02
2.8500E+00	-2.7788E-02	-3.1504E-02	-2.4356E-02
2.9000E+00	-2.7944E-02	-3.1681E-02	-2.4494E-02
2.9500E+00	-2.8099E-02	-3.1858E-02	-2.4631E-02
3.0000E+00	-2.8254E-02	-3.2035E-02	-2.4769E-02




Figure 129: AGP-8X IBIS Model Vs V-I Spec for Pullup and Pulldown

C. Sensitivity Analysis

C.1 AGP-4X

Several sensitivity analyses were performed to narrow the definition of the buffer and interconnect characteristics for 4X mode. These analyses used results from numerous simulations. Included in this appendix is the result of the final set of analyses that defined the buffer and interconnect characteristics. These results are displayed using 3D plots in Figure 130 through Figure 141.

These plots display the worst case skew at each point in the surface. These points are the same as the ones that are described in Section 2.2.3.4.

The first six plots are of the worst case for setup skew, and the last six plots are of the worst case results for hold skew. The buffer types are described in the figure titles.

Buffers labeled as “least linear” show some skew violations in these plots. These buffers also experienced numerous ringback violations, which are not shown.




Figure 130: Least Linear with a 1.4 V/ns Slew Rate




Figure 131: Least Linear with a 2.8 V/ns Slew Rate




Figure 132: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate




Figure 133: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate




Figure 134: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate




Figure 135: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate




Figure 136: Least Linear with a 1.4 V/ns Slew Rate




Figure 137: Least Linear with a 2.8 V/ns Slew Rate




Figure 138: 2.0 Curvature Compensation with a 1.4 V/ns Slew Rate




Figure 139: 2.0 Curvature Compensation with a 2.8 V/ns Slew Rate




Figure 140: 1.1 Curvature Compensation with a 1.4 V/ns Slew Rate




Figure 141: 1.1 Curvature Compensation with a 2.8 V/ns Slew Rate